



Science & Technology
Facilities Council

UK Research
and Innovation

Multi-VO Rucio and other Rucio work in the UK

Katy Ellis

On behalf of Andrew Lister, Ian Johnson, Teng Li and Alastair Dewhurst

11 Sept 2019, GDB at Fermilab


Overview

- UK support for non-LHC VOs
- Why Rucio
- Multi-VO Rucio
 - Development
 - Usage
 - To do
- Review of UK involvement with Rucio

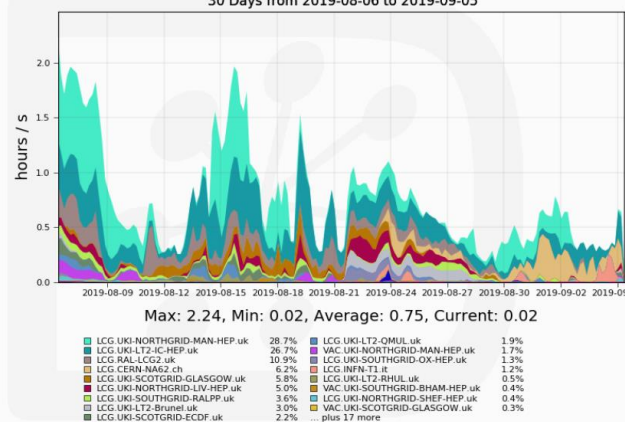
UK support for Non-LHC VOs

- The UK runs a **multi-VO DIRAC instance** to allow non-LHC VOs to utilize UK grid resources.
 - Used by T2K, SNO+, MICE, NA62, LSST, LZ ...
 - Support for individual users under the GridPP VO.
 - Allows future users such as SKA telescope, Diamond Light Source to evaluate services.

CPU usage by UserGroup
30 Days from 2019-08-06 to 2019-09-05



CPU usage by Site
30 Days from 2019-08-06 to 2019-09-05




Motivation for Rucio

- RAL provides a tape archival service for a large number of experiments/users.
 - Various data management tools are used including legacy software (e.g. LFC).
 - Individually they often don't have very much data (< 100TB) and very little effort.
- Rucio has the variety of features required to support these use cases.
 - Confident of long-term support
- The UK is looking to the future and making use of new storage technologies and APIs.
 - Rucio development/integration is more advanced than DIRAC.

Motivation for Rucio

- So for smaller VOs...

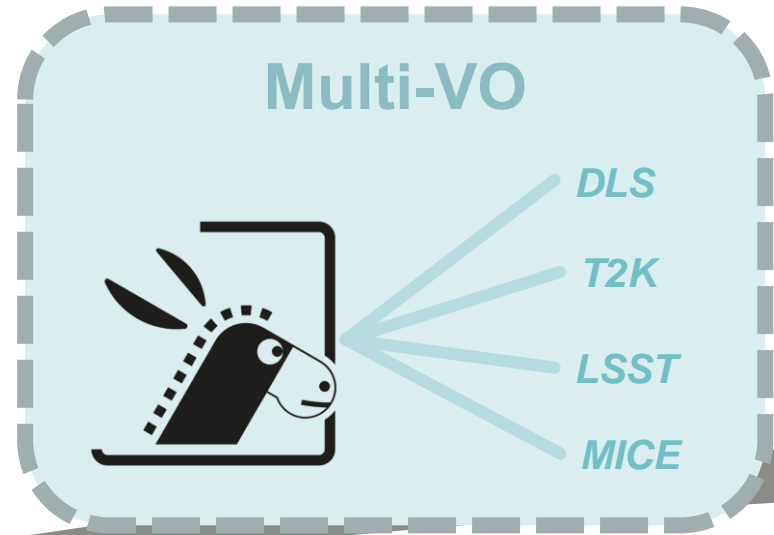



Continue to use DIRAC for workflow management
Use Rucio for data management



Multi-VO Rucio?

A single instance capable of supporting smaller VO's



Multi-VO?

- Offer as a service to any new experiments.
 - Can quickly set them and their storage endpoints up.
 - VO can focus on figuring out their computing model rather than setting up services.
- One instance to support and maintain.
- For large experiments, we will run a dedicated Rucio instance.
 - In future want to develop a way to split off instances.

Rucio Multi-VO Project

- Initially a 12 Month project at RAL to set up a multi-VO Rucio instance.
 - In collaboration with the Rucio team.
 - Mustn't have any impact to Single-VO instances.
- Bulk of development work done by Andrew Lister in a 6-month period ending September 2019.
 - Approximately 8000 lines of code changed.

Status

- Changes were split into 4 main pull requests:
 1. Separate internal and external representations for RSE
 2. Internal representations for Account and Scope
 3. Database Schema updates
 4. Multi-VO features!
- Most requests have been merged into the mainline Rucio code, one still pending.
- Two Rucio daemons still to address:
 - Auditor, which identifies dark data.
 - C3PO, which replicates hot files.

Usage (Admin)



super_root

Super_root creates VOs

root accounts

New VOs have individual root accounts for standard admin (e.g. Adding users)

users

Usage (User – Current)

An example of using Rucio for a typical (single-VO) user

Authentication

- Add Account, and credentials to “rucio.cfg” config file.

List Files

```
rucio ls user.jdoe:*
```

- Use Rucio CLI

Upload Files

```
rucio upload --pfn ~/a_file.txt --rse endpoint1 ...
```

- Use Rucio CLI

Usage (User – Multi-VO)

An example of using Multi-VO Rucio for a typical user

Authentication

- Add Account, **VO**, and credentials to “rucio.cfg” config file.

List Files

```
rucio ls user.jdoe:*
```

- Use Rucio CLI

Upload Files

```
rucio upload --pfn ~/a_file.txt --rse endpoint1 ...
```

- Use Rucio CLI

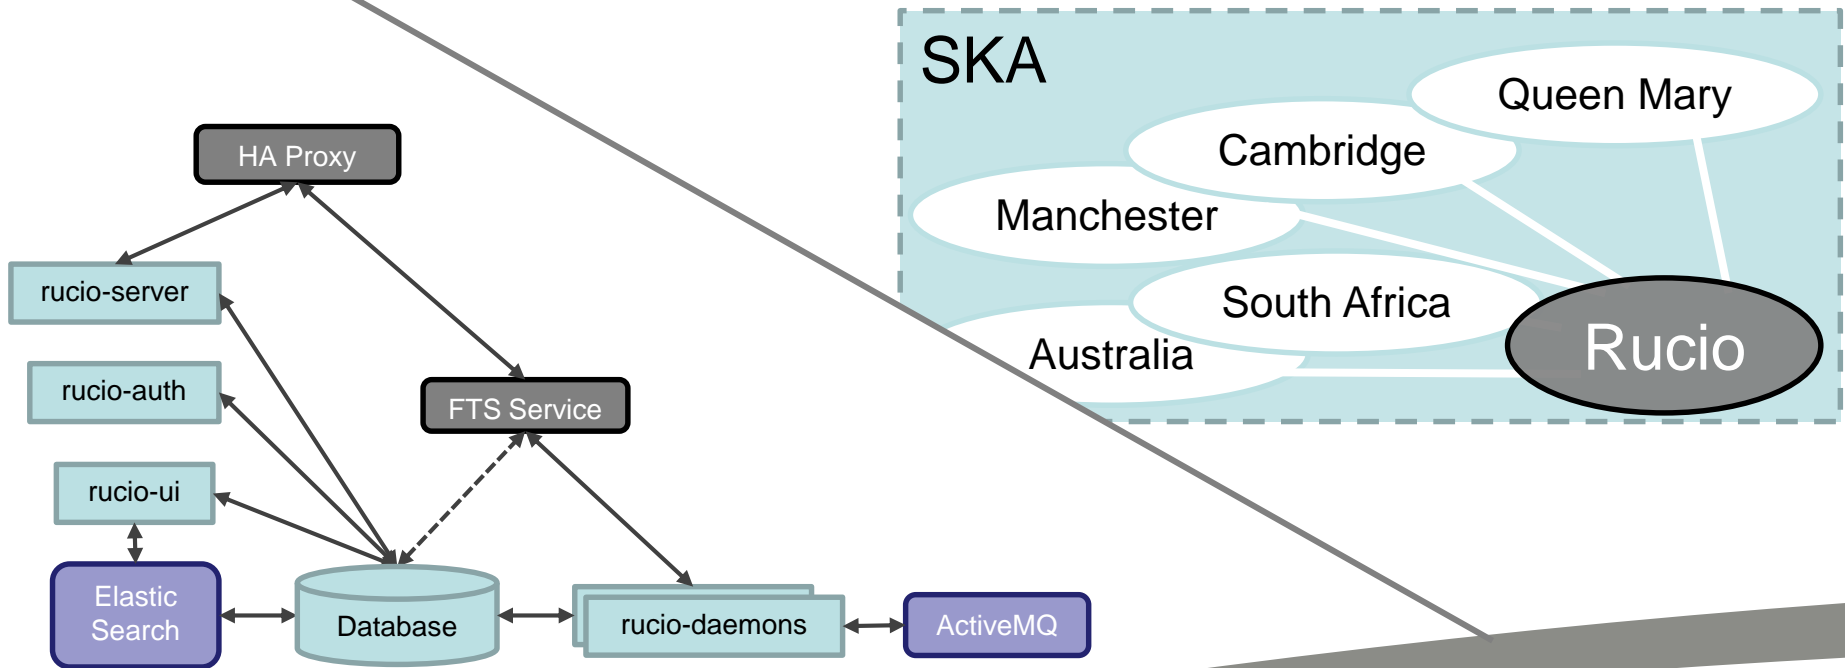
Next steps

- Adapt two remaining daemons.
- Demo Multi-VO at RAL for integration tests.
- Provide development instance for Multi-VO DIRAC integration.
- Production Multi-VO Rucio service running at RAL
 - Iterate permissions scheme as users join

UK efforts

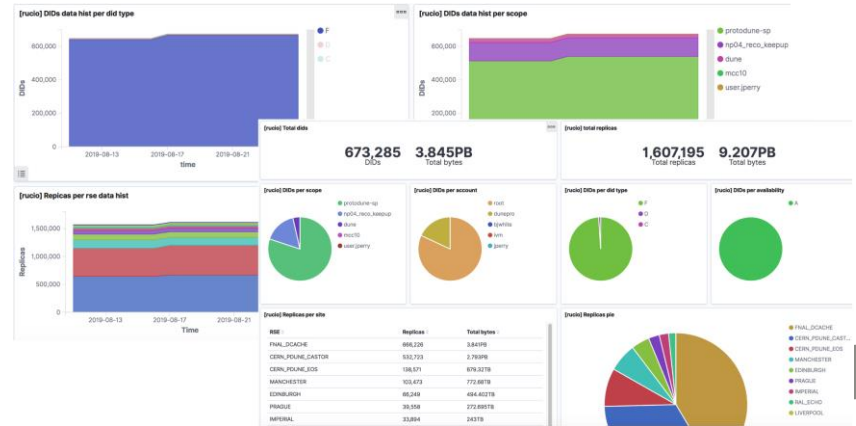
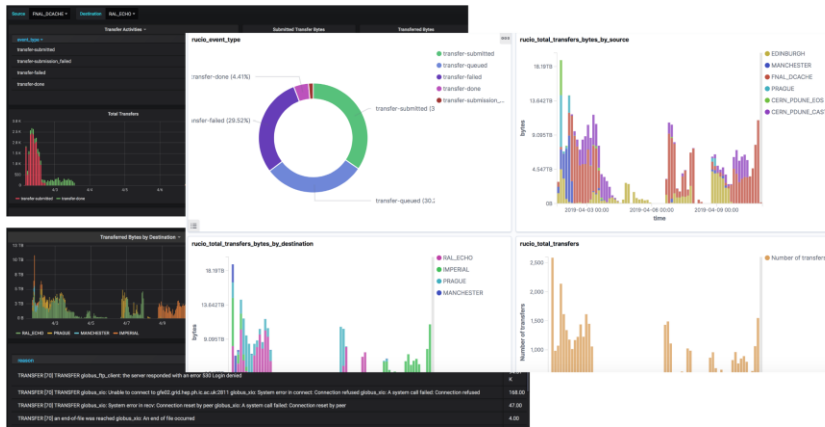
- RAL
 - Ian Johnson - Rucio@RAL service manager.
 - Katy Ellis - Working on getting Tape endpoints to work for CMS.
 - Eli Chadwick -Graduate starting Jan 2020.
- Edinburgh
 - Teng Li - Working on Rucio monitoring for DUNE.
 - James Perry - Working on Object Store integration for DUNE.
- Imperial College
 - Janusz Martyniak - Starting on DIRAC / Rucio integration

Rucio@RAL



Monitoring

- Teng has been working on Monitoring for DUNE.
 - <https://github.com/feipengsy/rucio-monitoring>
 - Will be added to Rucio@RAL



Rucio for CMS / RAL

- CMS will be using Rucio from Run 3, replacing PhEDEx.
- Integration activities ongoing, for example:
 - Production/User job submission
 - Database and monitoring setup
 - Large-scale testing
- RAL will be replacing the Castor tape service by 2022.
 - Still evaluating possible solutions.
- Mutually beneficial to work on Rucio Tape integration for CMS.

CMS Tape Testing

- CERN Tape Archive (CTA) replaces CASTOR at CERN
- Able to transfer into and out of CTA using Rucio
 - Must be done via EOS
 - Have performed small tests (~10 TB) working towards a more substantial test (~200TB)
- Functional test transfers working into all Tier-1 tape endpoints.
 - Tape recalls working from RAL only currently.
 - Other Tier-1 endpoints being worked on now.

Summary

- The RAL Tier-1 is committed to providing a data management service using Rucio for the many non-LHC VOs it supports.
- Development for multi-VO instance of Rucio is virtually complete.
 - Integration with multi-VO DIRAC is starting.
- UK staff are developing significant Rucio knowledge by contributing to DUNE and CMS work.



Science & Technology
Facilities Council

UK Research and Innovation