

Multi-VO DIRAC with multi-VO Rucio

Andrew Lister - STFC

andrew.lister@stfc.ac.uk



Overview

- Background within UK (GridPP/IRIS)
- Rucio
 - Rucio setup in UK
 - Work on Multi-VO Rucio
- Multi-VO DIRAC
- Plans for integration
- Summary

Background

- GridPP** - A collaboration of 18 sites across the UK dedicated to provide resources in the particle physics community via a large-scale computing grid (with a focus on LHC).
Currently in the process on applying for funding from 2020 - 2024.
- IRIS** - IRIS is community initiative with the aim of developing available resources for use in various physics domains across all of STFC

RAL Tier-1

RAL hosts the UK's Tier-1 data centre and its strategy is to provide a comprehensive set of data services:

- Echo disk storage (Grid + Cloud APIs)
- Tape archive (Castor but evolving in future)
- Rucio Service (Data management)
- FTS (Data movement)
- DynaFed Service (Federation and Authz service)



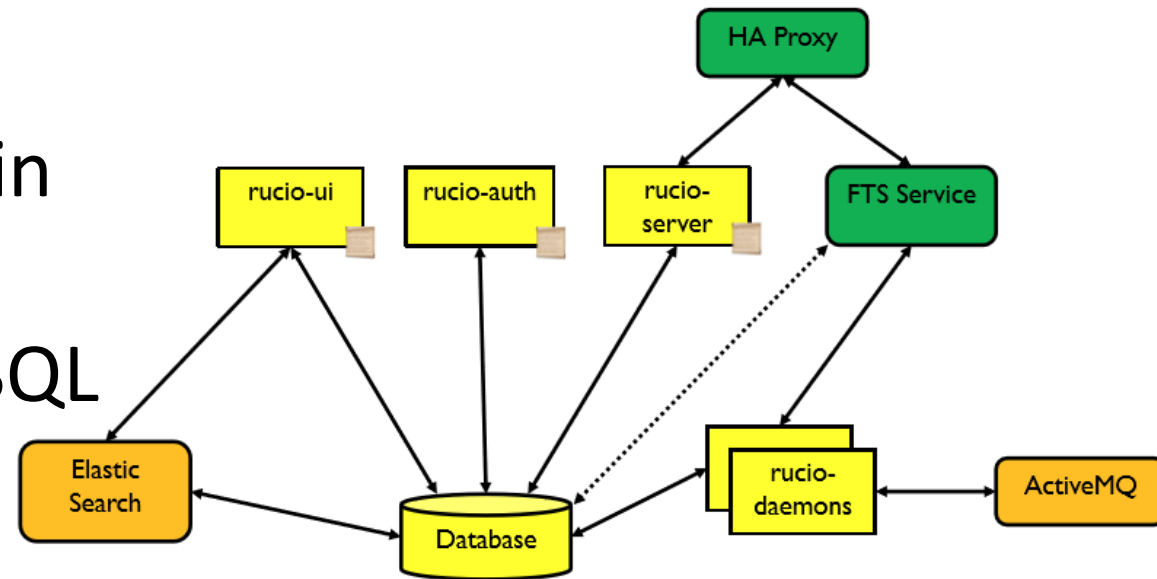
What is Rucio?

- Distributed data management tool
- Developed to handle HEP domain volumes of data
- Originally built for ATLAS, now supporting CMS and other VOs
- Continuation for lifetime of LHC (i.e. at least 15 years into the future)
- Open source

Rucio@RAL

Current instance set up for SKA VO,
managing transfers between UK sites and international
SKA partners in Australia and South Africa

Comprises a cluster of Rucio nodes in
OpenStack VMs and database
deployed on physical PostgreSQL
database server



Motivation for a Multi-VO

Small experiments:

- Often struggle to deal with different storage endpoints.
 - Abstraction layer between the user and underlying storage.
- Lack manpower for long term data assurance.
 - Rucio has a large community and long term support.

Rucio is deployed for large experiments currently

- Significant initial overhead puts people off.

Multi-VO DIRAC model at Imperial works well!

Multi-VO Rucio Requirements

Same as single-VO but with extra security

- Different VOs shouldn't see each others files
- Namespace must be able to overlap (e.g. same person working on two VOs)
- Users should be aware which VO they're working on

Mustn't impact existing instances.

What I'm doing!

Since starting this project:

- Travelled to collaborate with the core development team at CERN for 2 weeks.
- Created a plan for the Multi-VO extension
- Started development work and added code to the source to prepare for Multi-VO

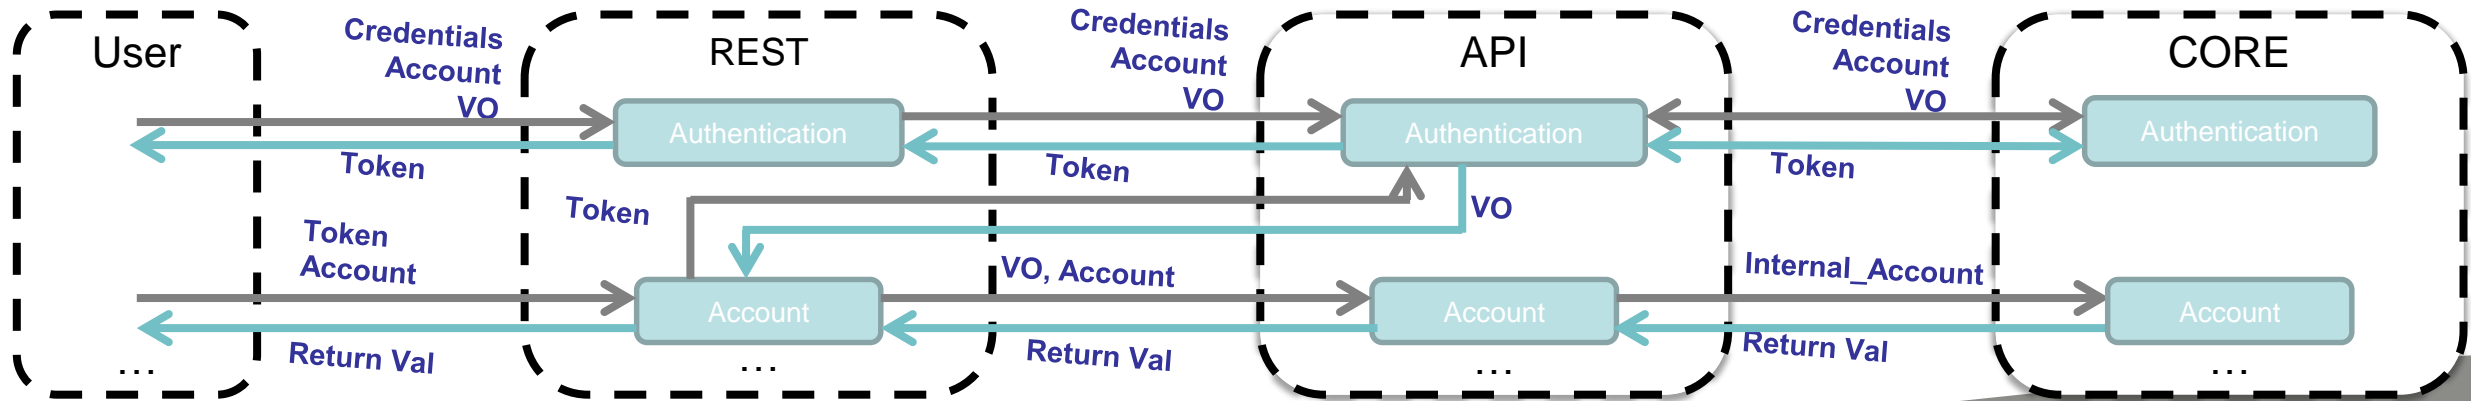
Still to do:

- Change the backend database
- Add VO variables throughout the code
- Create VO API for admins

Technical Challenges

We can't change the schema where existing instances have million row tables as they're physically ordered in memory

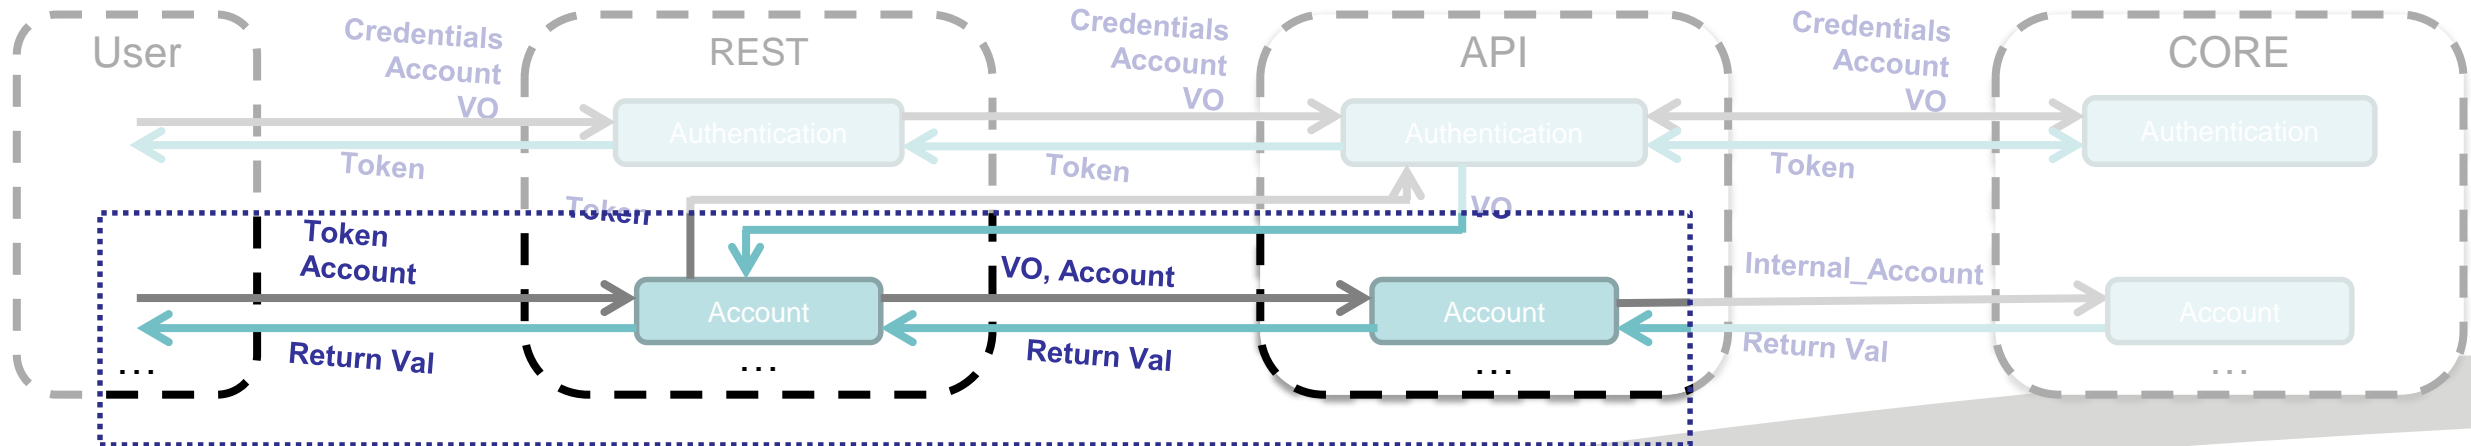
- Encode the VO into the Account and Scope identifiers



Technical Challenges

We can't change the schema where existing instances have million row tables as they're physically ordered in memory


- Encode the VO into the Account and Scope identifiers



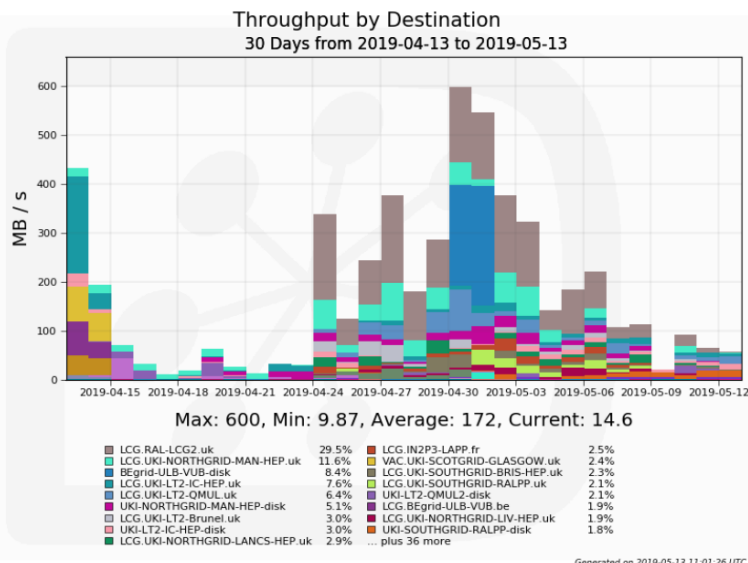
Multi-VO DIRAC

GridPP has a Multi-VO DIRAC hosted at Imperial since 2015.

- Successfully onboarding new experiments and allowing them to run at Scale.
- Limited support for non-Grid storage endpoints



– GridPP IRIS have chosen Rucio to manage storage endpoints.



Plans for integration

- Development of Multi-VO Rucio to finish by September.
 - In production by end of year.
- Janusz Martyniak has some effort to work on DIRAC side.
- Belle II (who you have just heard) are looking to integrate their Rucio instance with their DIRAC instance.

Summary

- GridPP have chosen to use Rucio for it's data management.
 - It would be great for everyone if both Rucio and DIRAC worked together!
- Multi-VO Rucio is scheduled for this year.
- DIRAC to start work on integration from September.