

Multi-VO Rucio

Timothy Noble - Rucio Service Owner
Rucio Workshop 10th November 2022

Multi-VO Rucio at RAL

- Rucio was initially setup in April 2018 at RAL to support SKA development
- A number of Graduates at RAL have worked with core dev team to develop Multi-VO functionality to support integration with Multi-VO DIRAC to support experiments
- Evolved to supporting Multi-VO
- Has longer-term support with 1 FTE at RAL, as well as a Graduate
- George Matthews, Started in September working on Rucio until March

RAL Rucio Funding

Current Funding

- EGI funding for running Rucio to support VOs interest in using Multi-VO Rucio. Currently, no direct uptake instead opting for other solutions or deploying their own instance.
- Swift-HEP funding for research – The instance being used to investigate the best was forward for UK dataflows for HL-LHC

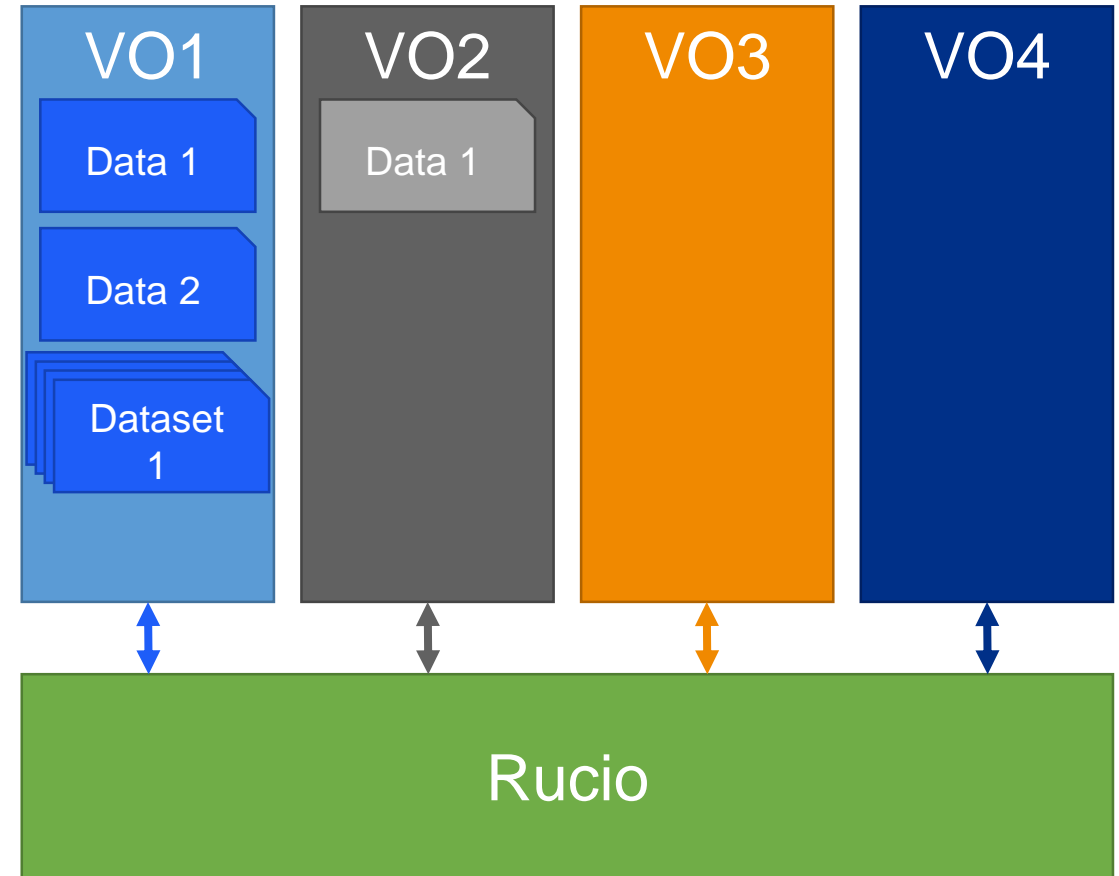
Future Funding

- SKA
- LSST

Separation of VOs

VO namespaces

- Each VO is a separate namespace
- Policy packages can be VO specific
 - Placement algorithms
 - Authentication levels
 - Regex of for sense checking



Advantages to use Multi-VO Rucio

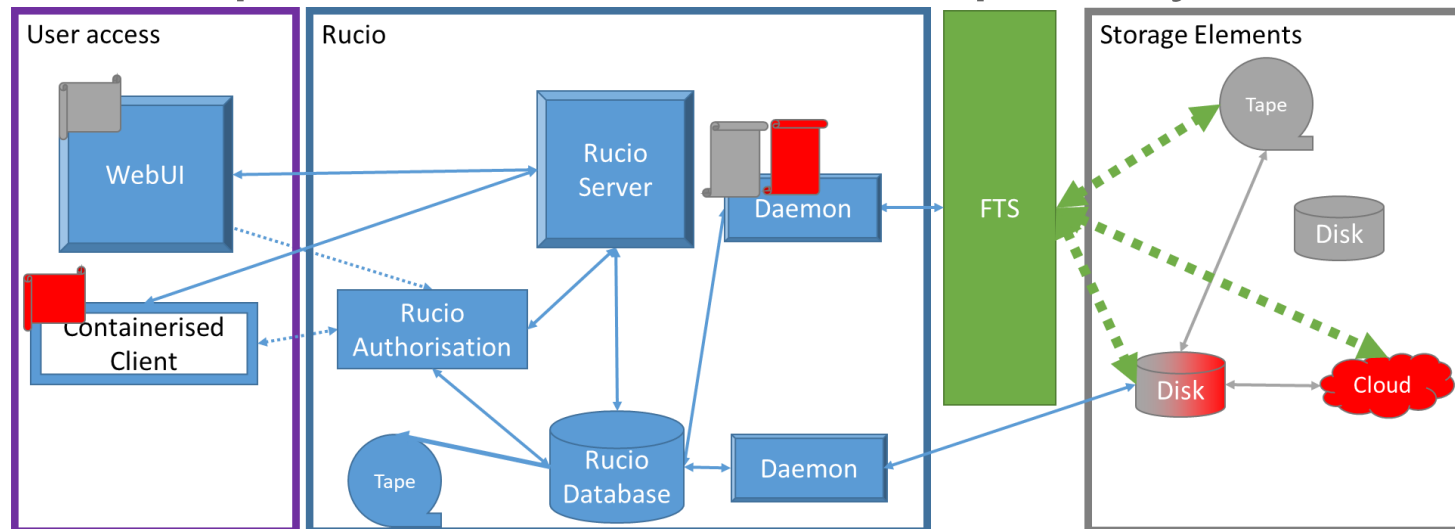
- Running a Rucio instance that supports multiple VOs is beneficial for small experiments:
 - Maintained by RAL not by smaller experiments
 - Low levels of load from smaller experiments.
 - One instance to support and maintain.
 - Shared RSE configuration
 - New VOs are quick to add – work with VO admin to setup their environment
- More contact with Developers and larger communities using Rucio to know how best to utilise

Multi-VO Setup

- Configuration switch to enable
 - Added VO argument to client to specify (can be transparent to users if set in ENV or rucio.cfg)
 - Addition of credentials (x509 proxy)
- Originally pip installed Rucio server and daemons
- Deployed components to Docker
- Daemons and server deployed on K8S cluster
- Currently supporting 2 VOs simultaneously with single Daemons

Feature development: Multi-VO Conveyor

- Multi-VO Rucio was still under development when I joined
- Conveyor daemons could only have once authentication x509
- Certificate selection based on VO for single daemon operation
 - Implemented in 1.27
- Completed the Multi-VO capability for the daemons



Multi-VO deployment developments

- Transition to a K8S deployment of Rucio
Convert cronjobs for testing over to K8S jobs
- Adapt existing Rucio cronjobs for multi-VO K8S deployment
- Deployment of a small Kubernetes cluster to support Rucio
 - Rancher RKE2
 - Helm
 - Prometheus monitoring from Graphite
 - Crons – renewals and delegation
 - Server, daemons
 - TODO: Flux, authentication server
 - External deployments
 - Loadbalancer / HAproxy

Graduate Project


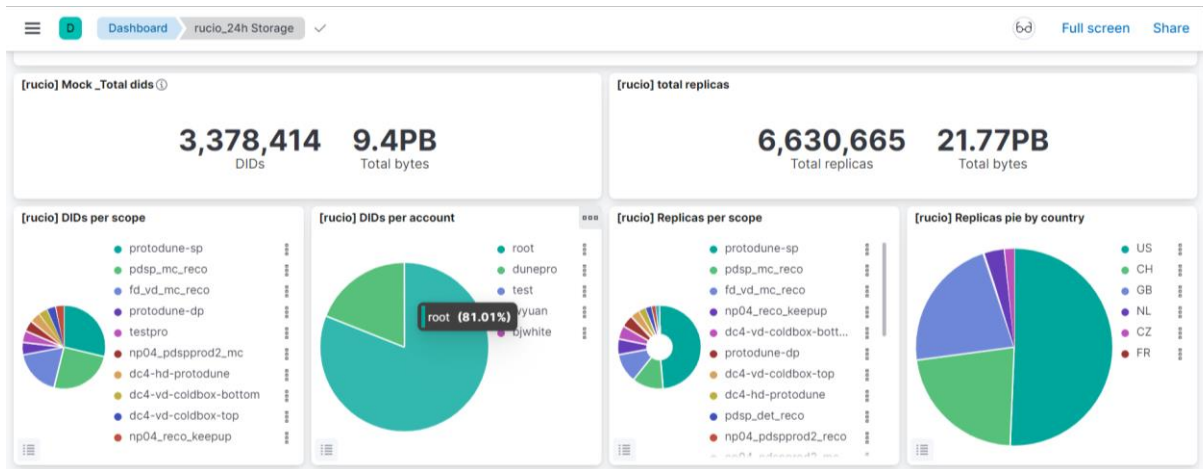
- George is to be working on Rucio for 6 months
- Started in September
- Working with Mayank and other Devs to ensure new WebUI will function with and has features specific to Multi-VO
- Helping deploy functionality for the K8S
 - Converted testing crons to K8S deployment
 - Put in PR to convert proxy renewal and delegation to K8S proxy

Token Authentication

- In principle works on Multi-VO, needs some configuration and editing of sync scripts
- Transition to K8S deployment 1st step
- Information for IRIS-HEP and EGI Check-in OIDC is already configured and ready for testing as soon as auth server is ready and functional

Rucio Monitoring

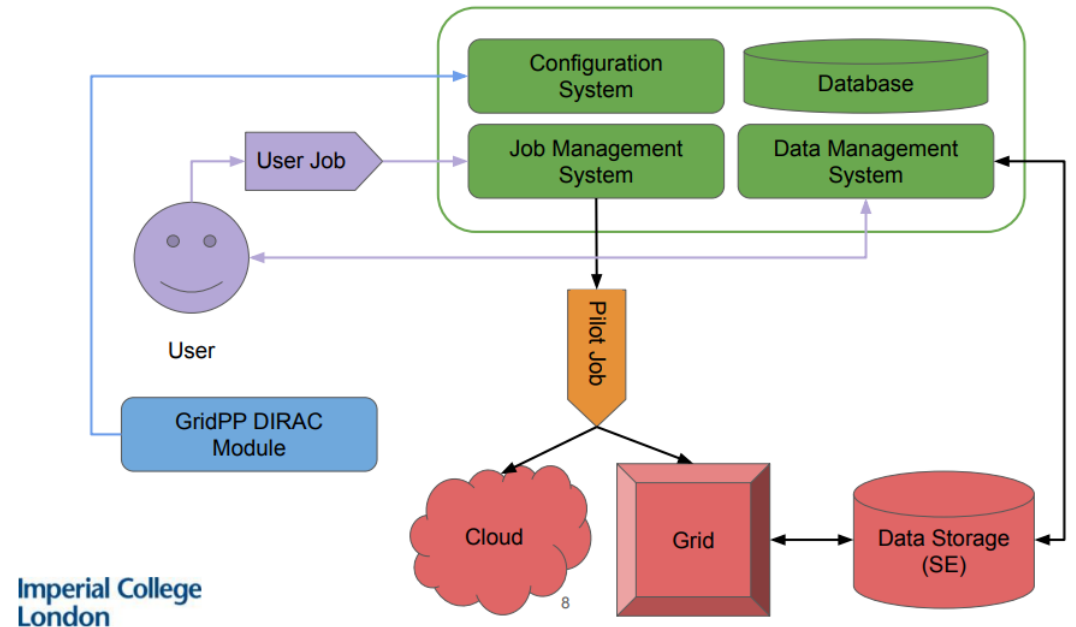
- Implementing further monitoring for Rucio by collaborating with Edinburgh (Wenlong and Rob) to get the Hermes-based monitoring implemented and accessible by communities that may want to use Multi-VO Rucio.
- Modelled on the Monitoring deployed for DUNE



Data-lake within workflow management

- Requirement to move to Rucio 1.29
- VO specific policies to ensure correct working with DIRAC without disrupting other VOs on Multi-VO Rucio
- Work with Janusz to ensure integration with DIRAC

DIRAC Schematic (for reference)



Janusz's presentation at [SWIFT-HEP May meeting](#)

Rucio development work

- Core Rucio work:
 - Removing deprecated storage code (S3 related) to cut down dependencies
 - Debugging and fixing a problem with S3 RSEs in Multi VO mode
 - Enforcing version check on policy packages to prevent breakages when Rucio API changes
- Improvements to DUNE policy package:
 - Updating to work with newer Rucio versions
 - Using MetaCat Python API instead of direct REST calls (with help from Igor Mandrichenko)
 - New SURL algorithm that queries MetaCat, for use with tape storage sites

Questions?