


# Evaluating Rucio for SKA Regional Centre Data Management

AENEAS : An SKA Regional Centre for Europe Develop a concept and design for a distributed, federated European Science Data Centre (ESDC) to support the astronomical community in achieving the scientific goals of the Square Kilometre Array



Credit: SRAO  
Heywood et al,  
2019

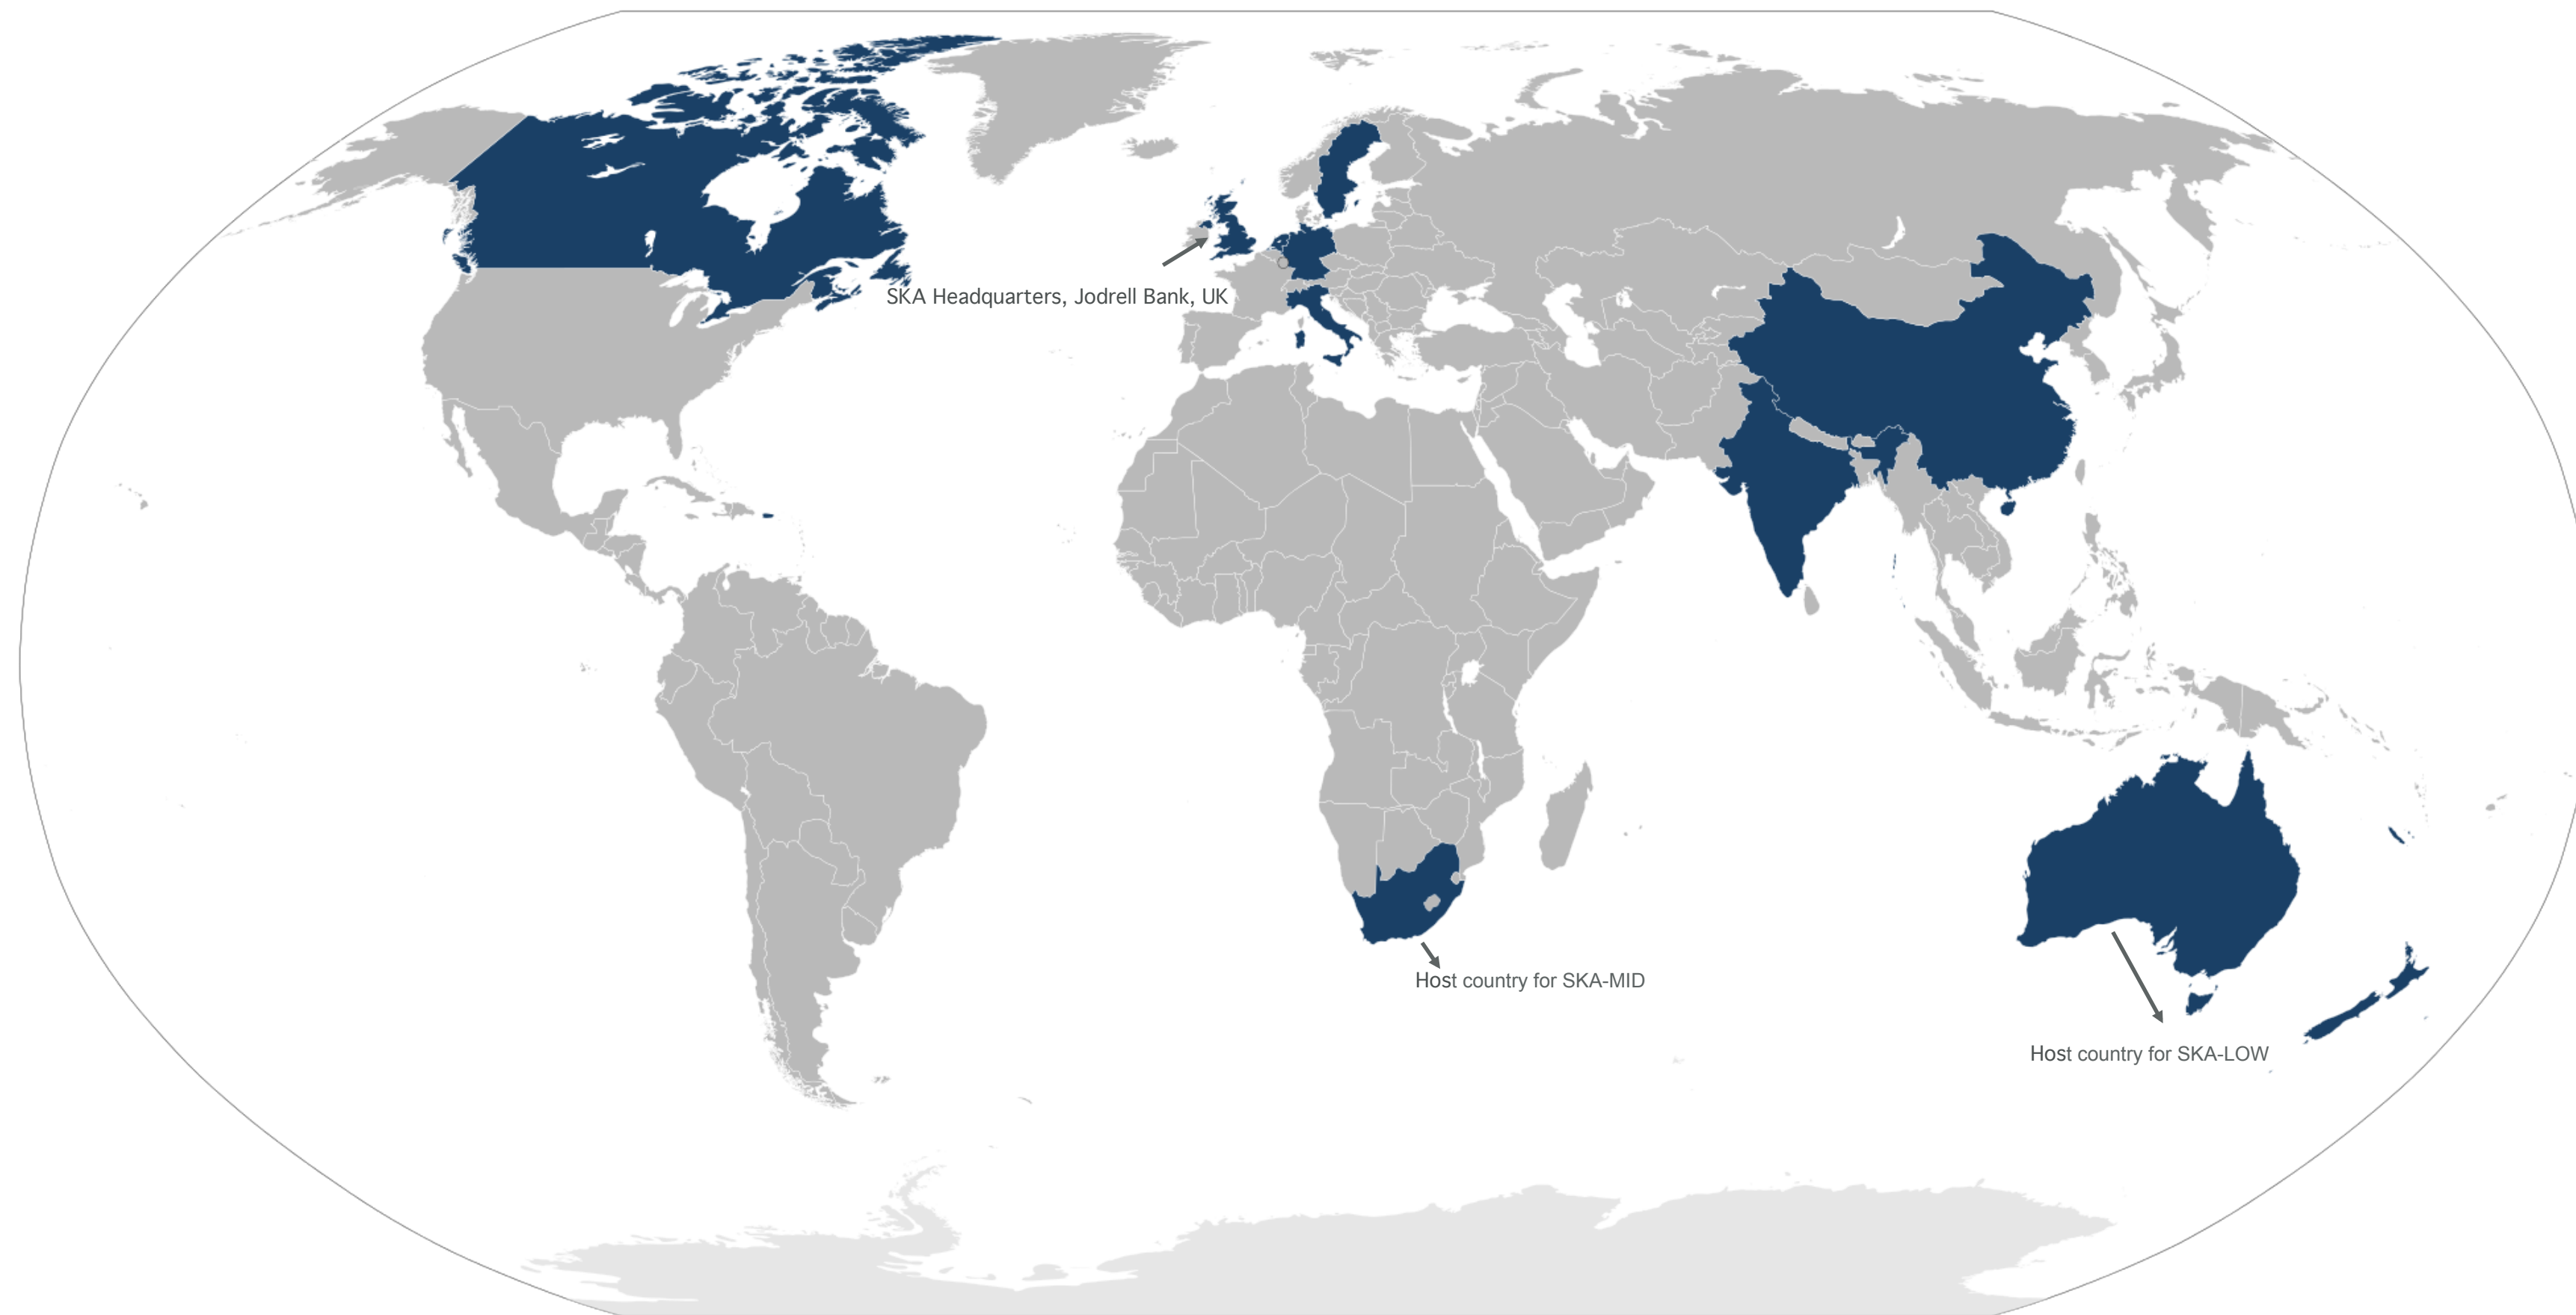
# THE SQUARE KILOMETRE ARRAY

Science Data Processors at both the SKA sites will collect, process and churn out science data products that will be pushed out to one or more regional centres around the world

## MEMBER COUNTRIES

- Australia
- Canada
- China
- India
- Italy
- Netherlands
- South Africa
- Sweden
- France
- Germany
- Spain
- UK


Potential new members:  
Portugal, Switzerland, others





# SKA REGIONAL CENTRES

SKA Regional centres will provide a platform for data access, data distribution, post-processing, archival storage, and software development.




## ARCHIVE

Archival of the raw data but also once scientific results are published, but outputs of analysis are made available.



## DATA DISCOVERY

SDP has pushed the data to the regional centres, now what? How will users find/peruse their data. How will published results be easily found?




## DISTRIBUTED DATA PROCESSING

Use cases are made to be reproducible. Compute comes to the data (high data volume).



## USER SUPPORT

SRCs must support the key science project teams as well as general users. This will mean user ability with be varied.



## INTEROPERABILITY

Multiple regional SRCs, locally resourced but interoperable. SRCs may be heterogeneous in nature but with common core functionality

# TYPICAL USE CASES

## ➤ CALIBRATION OF UV DATA

Direction independent calibration is a multi stage process with alternate parallel and aggregation steps. Input data ~ 3-4 TB reduced to 500 – 700 GB. Direction dependent calibration is a lot more computationally/memory intense , ~256

## ➤ <sup>GB RAM</sup> PULSAR RE-FOLDING (GENERATING TEST VECTORS)

Single core, Low memory jobs that produce outputs of the order few 10s of GB. Simulating test vectors single core, 30-40 GB RAM, produces outputs with similar size

## ➤ IMAGING OF CALIBRATED DATA

Highly computationally intensive, memory intensive due to nature of algorithm

## ➤ MACHINE LEARNING

Convolutional Neural Network to detect and classify large number of galaxies found in an image (GPU required for training)

## ➤ <sup>Random Forest classification for multi-wavelength photometry and finding new sources</sup> DATA PRODUCTS

These will vary from few GB per observation to several hundred. And data archives over time vary from few hundred GB to few hundred PB (from experiment to experiment). Users consuming data will also be generating secondary data which may not be smaller than raw data

# WHAT is RUCIO

## ➤ DISTRIBUTED DATA MANAGEMENT SERVICE

Manage user data in a federated way across many physical sites. Integration of scientific and commercial storage.

## ➤ PARAMETERISED REPLICATION/DELETION OF DATA

User friendly ways to requesting data replication/deletion allows for orchestrated data flow. Set rules like “I want 1 copy of this dataset on disk at Manchester for 1 year and 1 copy on tape in South Africa always”

## ➤ DEMONSTRATED TO WORK AT SCALE

Developed for and by the ATLAS project, with more communities adopting it.

Shown to handle 1B+ files, 450 PB+ data, across 120 data centres, 5 HPCs, 2 Clouds, 1000 users

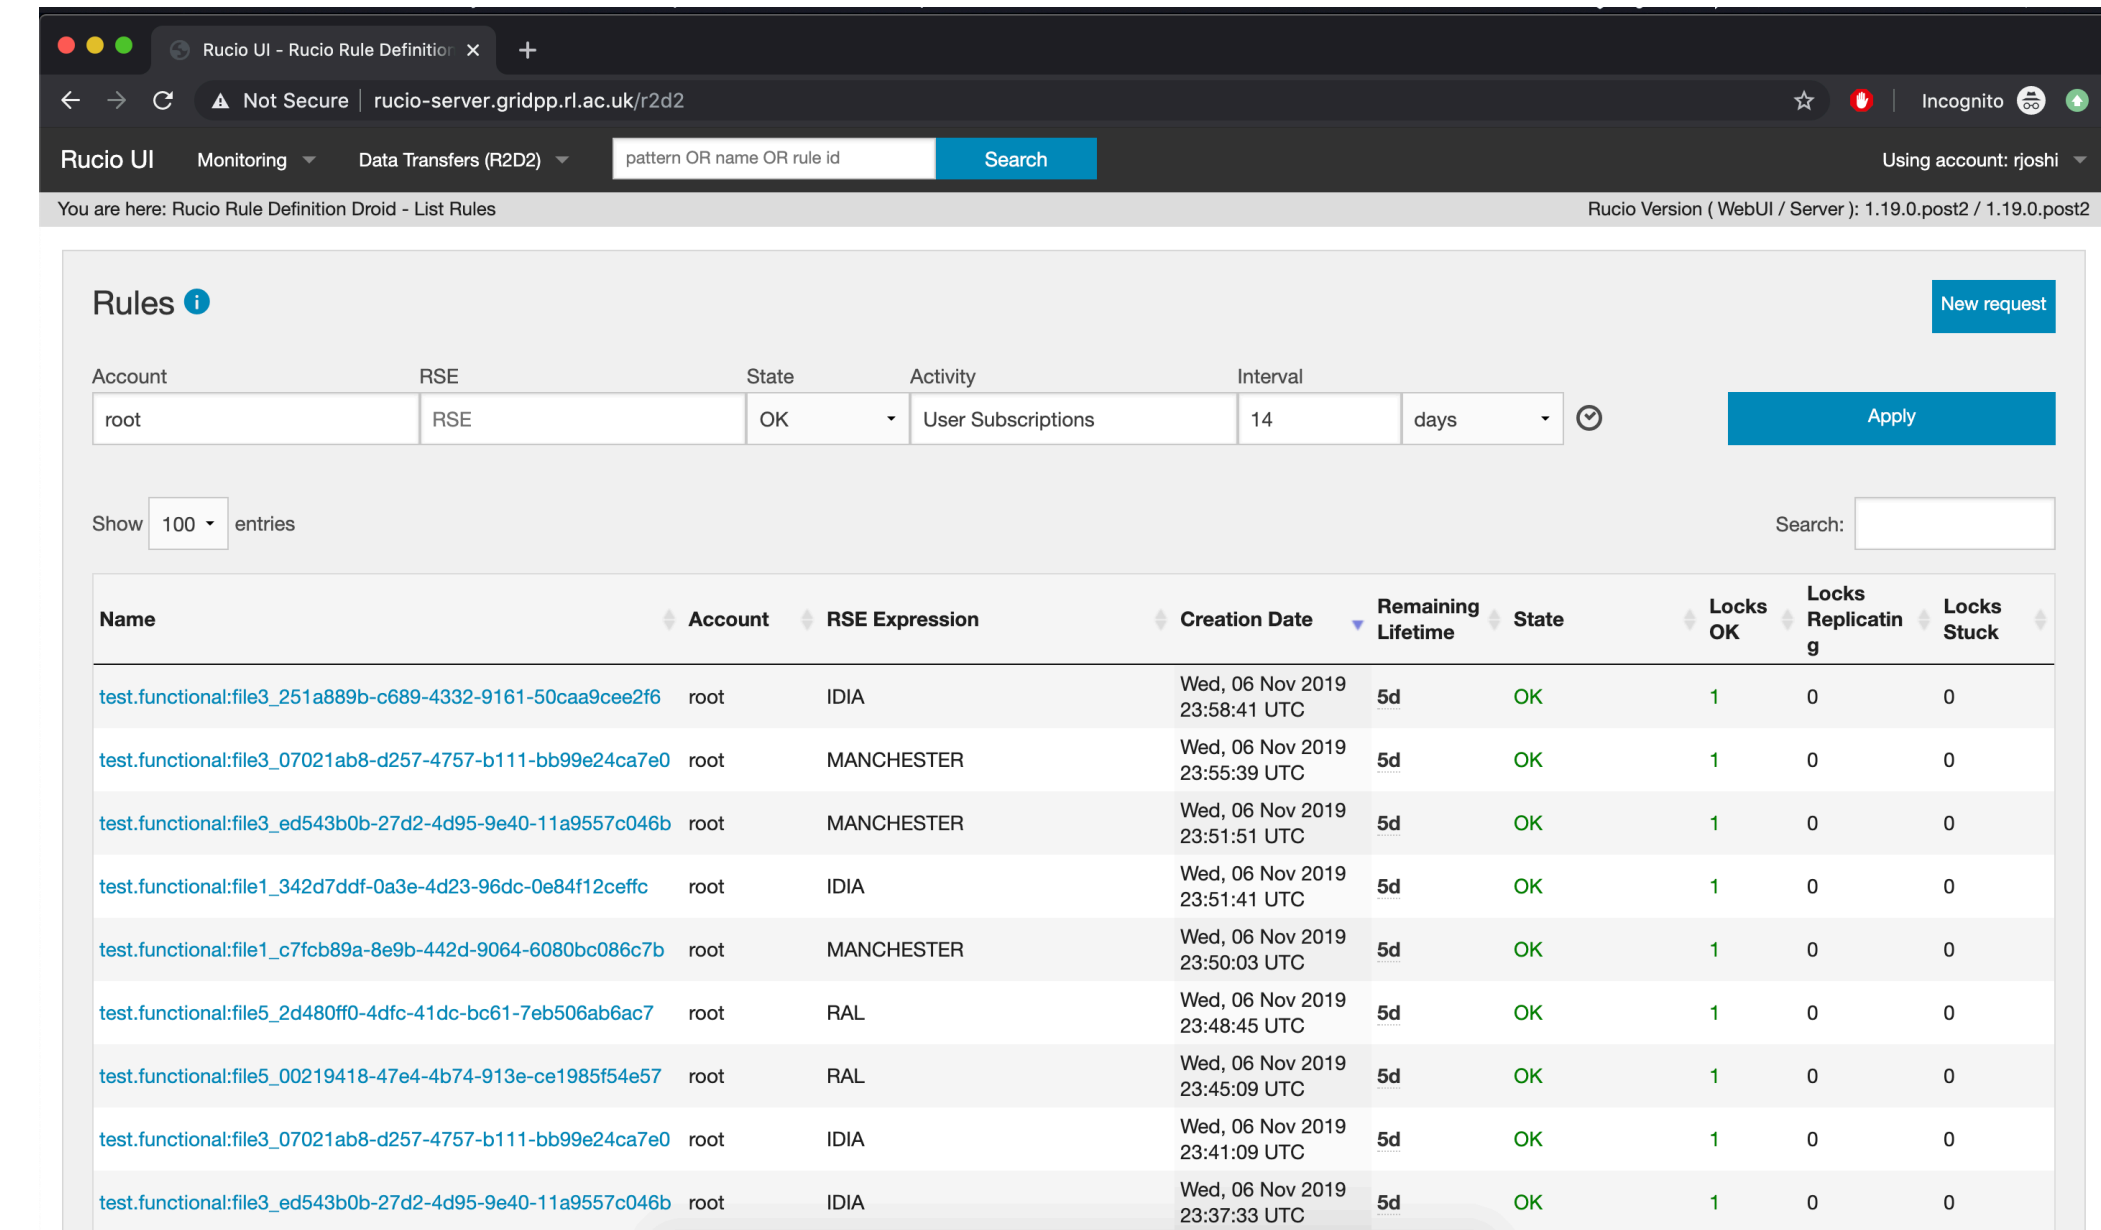
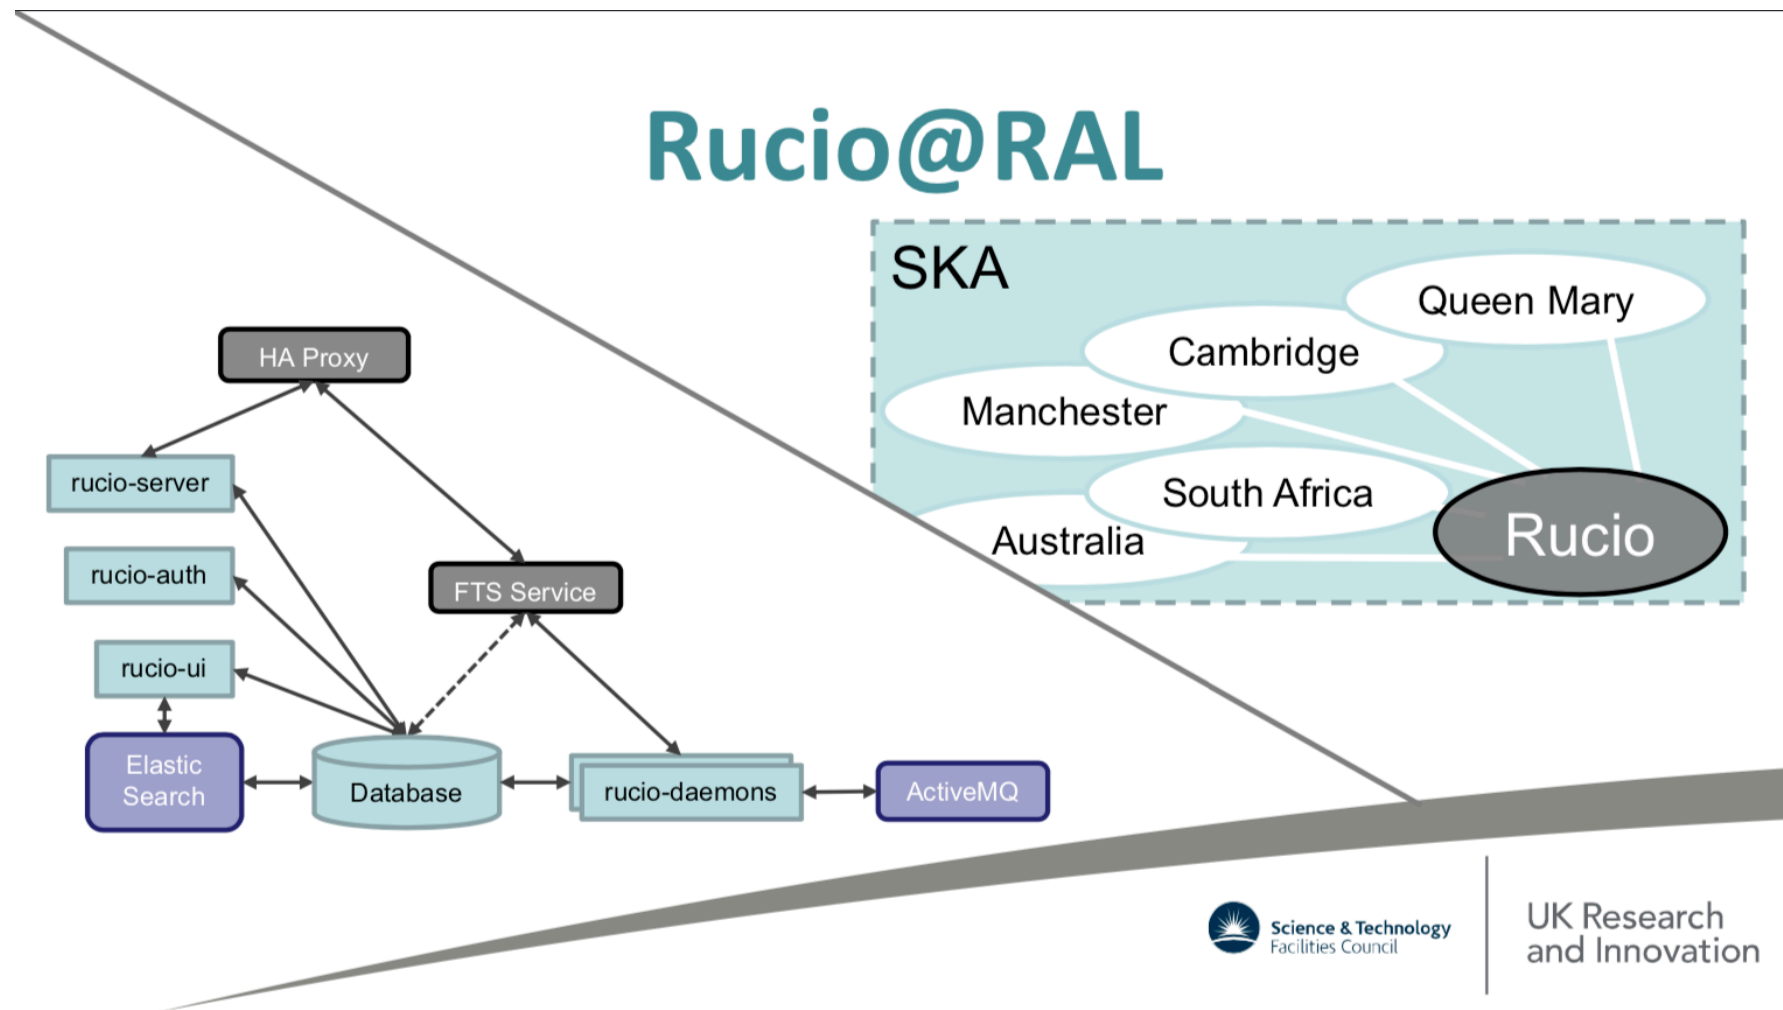
## ➤ RUCIO FUNCTIONALITY

Web-UI, CLI, and API to discover/download/upload/transfer/annotate data. It is currently able to use PANDA as a workload management service but APIs allow for integration with other WMS as well (for eg – DIRAC-Rucio integration

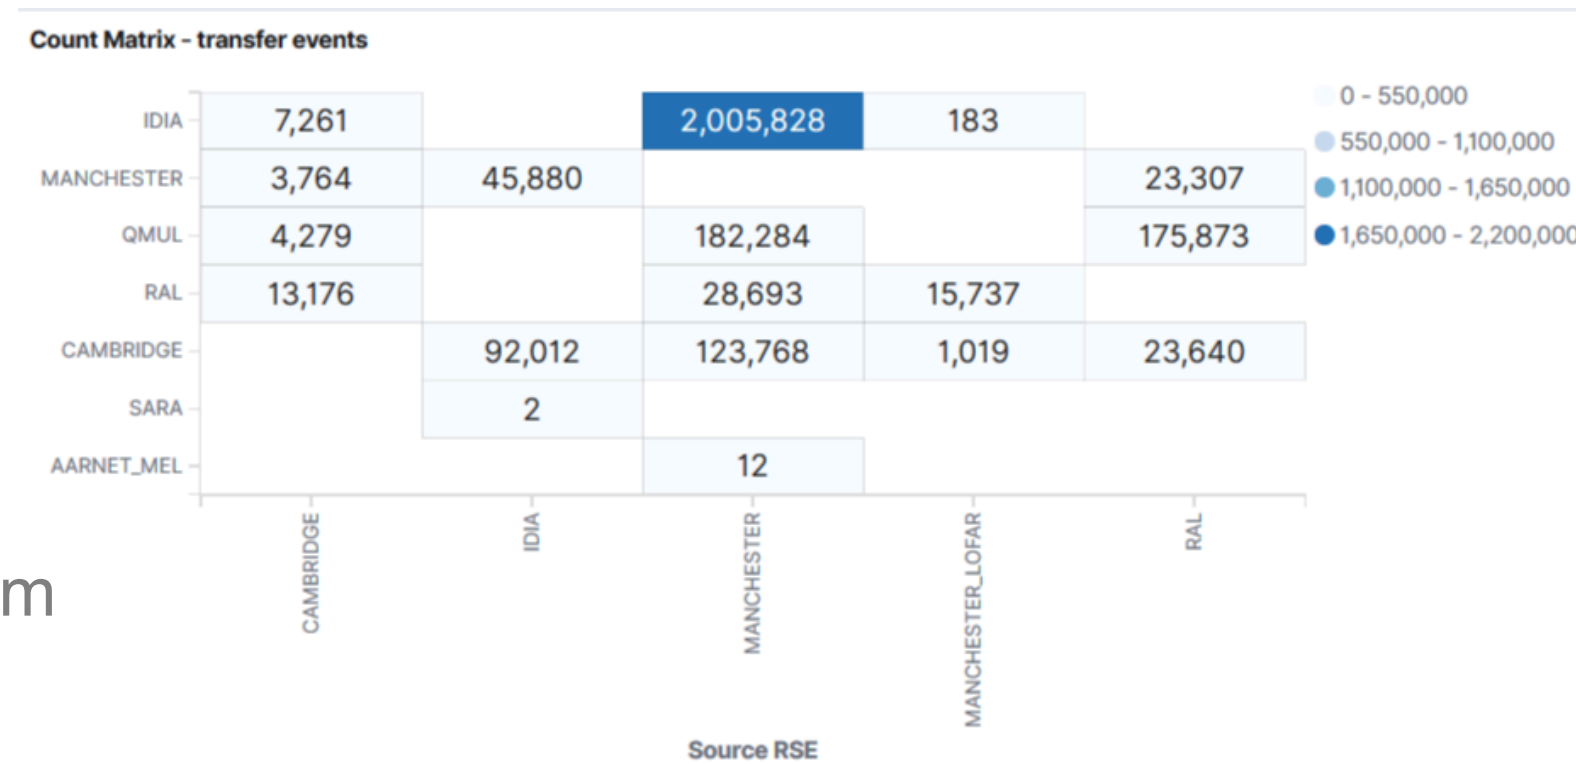
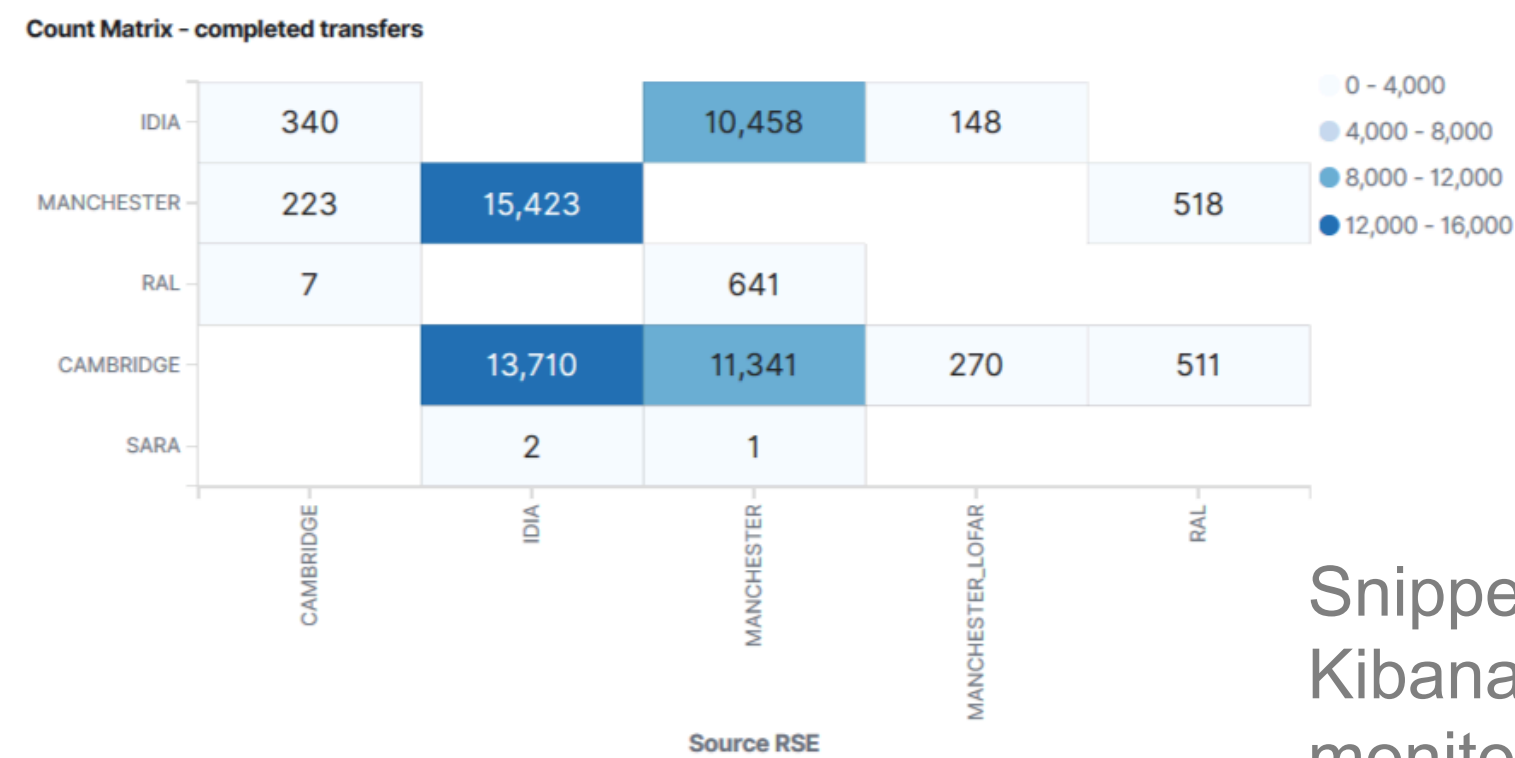
## ➤ OPEN SOFTWARE

Rucio is an open-source software licenced under Apache v2.0. Makes use of open source toolchains. Aims to adhere to FAIR principles

# RUCIO at RAL



Thanks to Andrew Lister for the slide



Snippets from Kibana monitoring dashboard

Thank you to folks at RAL for setting up and maintaining the instance – Alastair Dewhurst, Ian Johnson, Andrew Lister, Eli Chadwick



# WHAT WE HAVE TESTED SO FAR

Data being used is astronomy data from pathfinder instruments (LOFAR) although the volumes are not similar to SKA data volumes

## UPLOAD DATA

- Upload data using Rucio CLI
- Registering data with Rucio that is already uploaded to storage

## DATA REPLICATION


- Replicate data using Rucio CLI
- Parameterised replication rules
- Moved data around sites in UK, SURFsara and IDIA
- Functional tests to demonstrate a full network mesh test

## DATA DELETION

- Delete data using Rucio CLI
- Data deletion due to rule expiration

## OTHER

- Alternative methods of authentication when using the CLI (other than x.509)
- Add new Rucio storage elements
- ELK stack for monitoring (Elasticsearch, Logstash, Kibana)



# CHALLENGES

( k n o w n )

DEVELOPING TRUST IN THE USER COMMUNITY TOWARDS THE SYSTEM



# LOOKING AHEAD

## ➤ METADATA BASED SEARCHING, PERMISSIONS

Yet to explore how well metadata based searching will scale. Since astronomers like to keep their data protected, at least within the embargo period, some development needed to add configurable usage policies.

## ➤ WMS INTEGRATION

Needed to show a functional use case end to end. DIRAC – Rucio integration is an ongoing effort by folks at Imperial, with interest to support from ESCAPE.

## ➤ ESCAPE WP2 DATA LAKE PROTOTYPE

Rucio is the primary candidate for the storage orchestration service in the data lake prototype. Work will continue in the form of user experience and development as part of the ESCAPE project.

## ➤ MORE NODES, IMPROVED THROUGHPUT

Setting up Australian RSE, look into the possibility to use Rucio to deliver MeerKAT/ASKAP data to Europe (this would require higher bandwidth). Do you have some storage you'd like to add to the Rucio testbed? Get in touch!