

Data Management Update

Timothy Noble - Rucio Service Owner
SWIFT-HEP/GridPP 24th March 2022

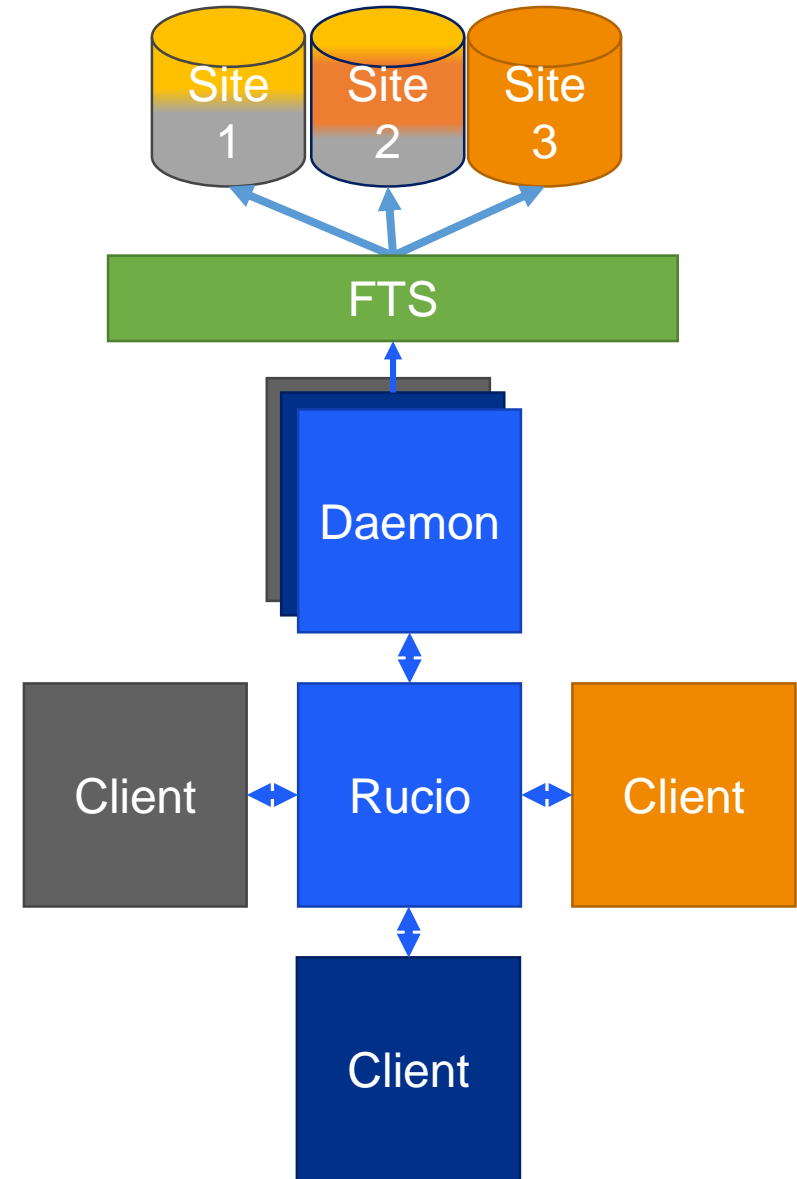
Work package 1

- Set up a UK data-lake prototype.
 - This will build on the DOMA prototype, with the intention that there will be one data lake per region/country.
 - The UK data-lake will use the Rucio instance that is already available at RAL (and being developed with IRIS funds).
- Setting up the data-lake in the first instance consists of 3 steps
 - Configure core sites
 - Configure additional – different storage sites
 - Generate metrics for comparison

Rucio



- Data management tool
 - Integrates with many storage solutions
 - Data can be stored across multiple sites, with different setups and protocols
 - Rules based fulfilment
 - Intelligent data movement to fulfil rules
- Close integration with FTS
- Open source
 - Community-driven development



Rucio internal developments

- Development of Multi-VO features to include conveyor daemons
 - Running since December with a single transfer daemon running 2 VOs
- Move from cloud VMs to containerised deployment of daemons
 - Looking to move to K8S deployment
 - Up to date deployment of Rucio
- Maintain Multi-VO Rucio to facilitate DIRAC integration compatibility

Rucio External Developments

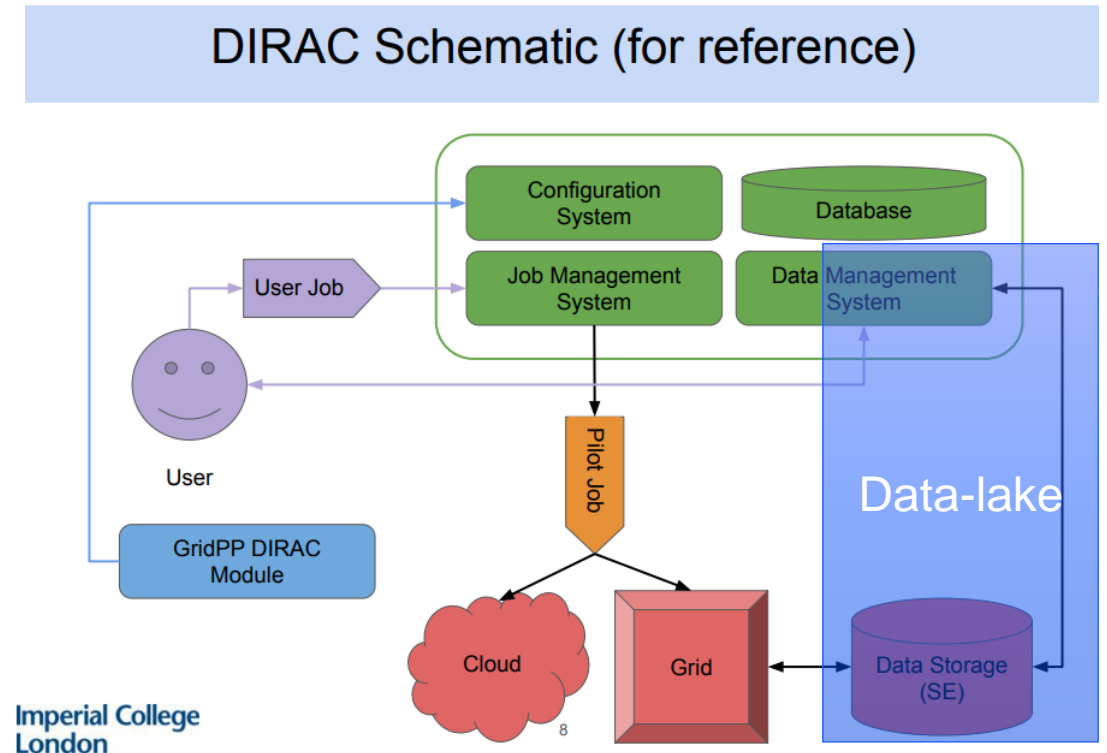
- Working towards integrating S3 endpoints
- Integration with IRIS-IAM
 - Configured IRIS-IAM client
 - Internal Rucio configuration issues,
 - K8S a step towards the answer?
 - If networking issues resolved
- Integration with EGI Check-in

Development of a prototype Data-lake

- Configuration of 4 core sites, preferably using different storage technologies
 - Currently Multi-VO Rucio is set up to use 9 functional storage endpoints in the UK and one from France. Of these endpoints the following storage technologies:
 - Ceph
 - dCache
 - DPM
 - StoRM
 - XCache - Edinburgh
 - And the following protocols are being used:
 - SRM
 - GSIFTP
 - WebDAV – all using this, with other protocols as alternatives

Data-lake within workflow management

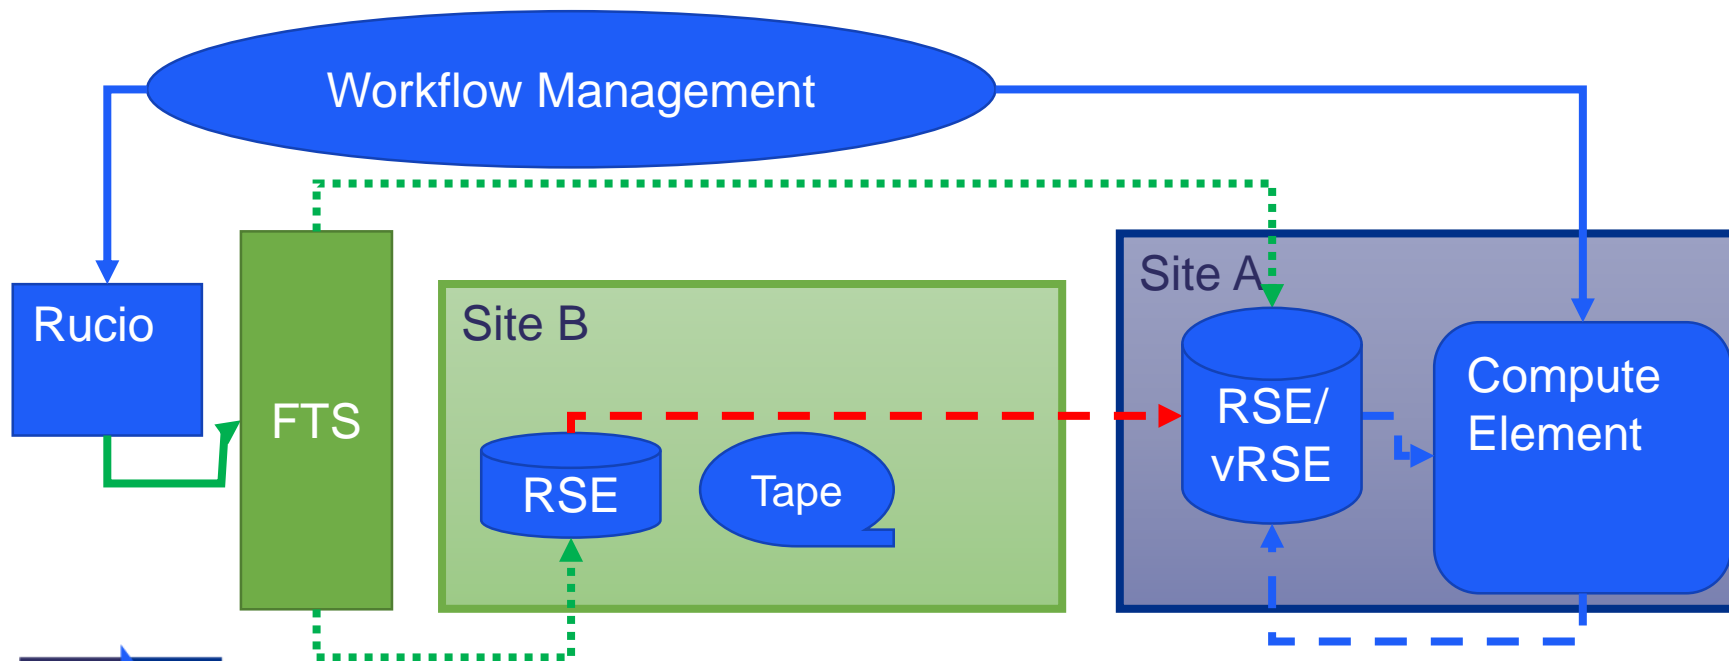
- Interested to see data movement for the completion of jobs
- Is the plan for the use of pre-placement in Caches or RSEs?



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

Data management to support jobs

- RSEs within Rucio come in 3 main types: Standard, Tape, Cache (vRSE)
- The job is registered at the compute element and ask Rucio to pre-place the data – e.g. Birmingham



Pre-placement using Rucio

RSE

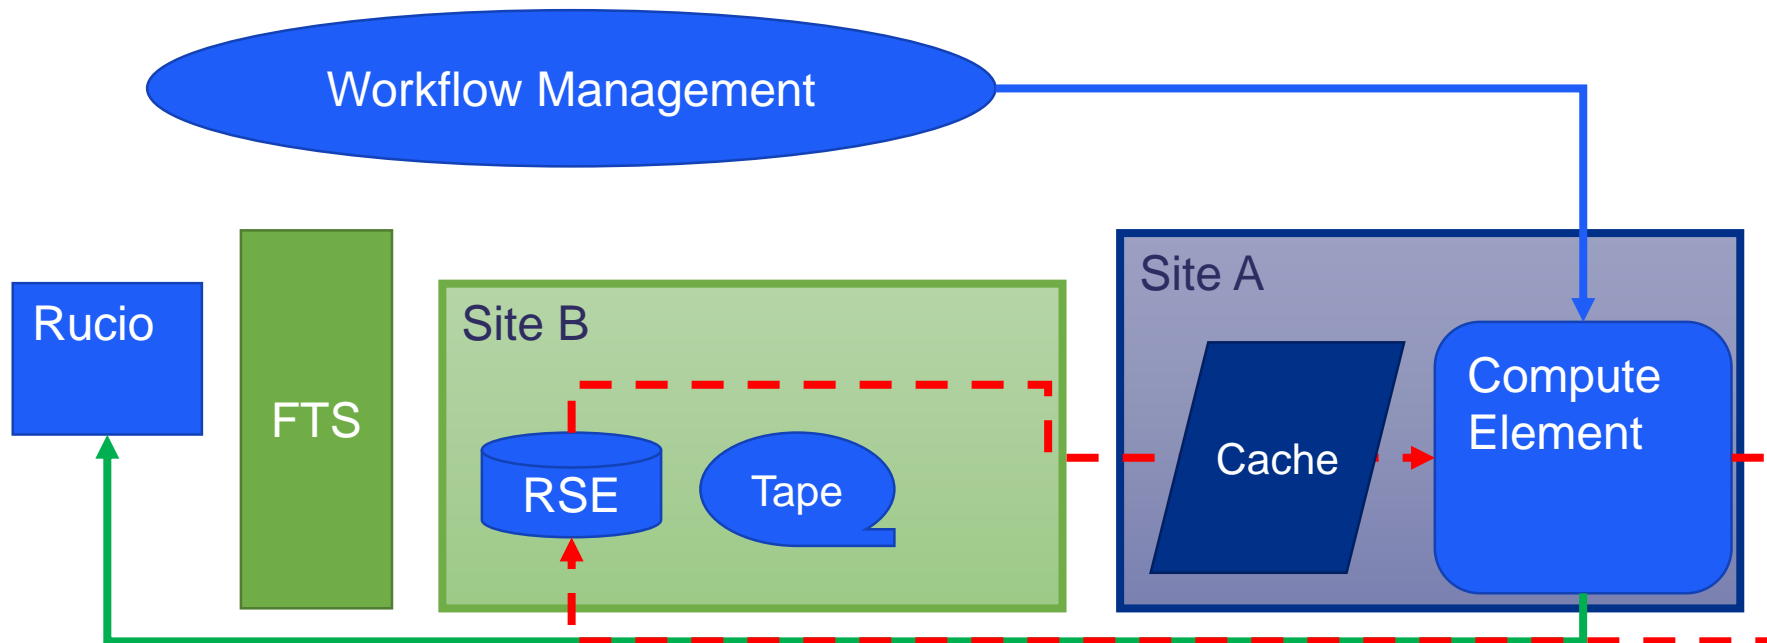
- Data is subject to Rucio's normal replication management

vRSE

- Data is excluded from Rucio's normal Rule maintenance
- Data moved here is managed by the workflow management tool
- Registering and deletion handled by workflow management

Data management outside of Rucio

- Reactive XCache e.g. Oxford

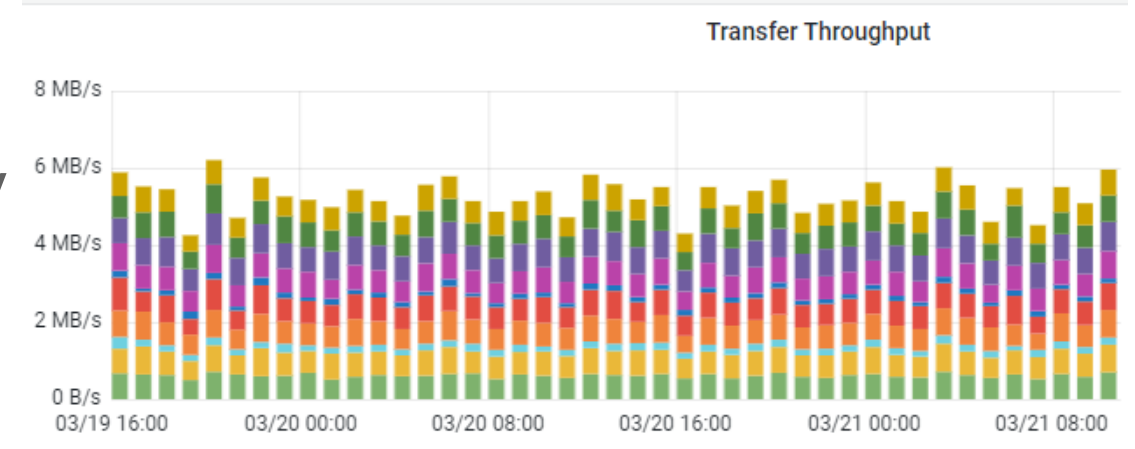



Expansion the beyond core sites

- Configure additional sites to access data. As well as the core sites, we will configure at least one site with state-less storage (i.e. a cache), one site with no local storage at all and one external site such as a Cloud Compute or an HPC to access the data-lake.
- Work from Birmingham showing XCache pre-placement working well
https://indico.cern.ch/event/1128343/contributions/4787170/attachments/2412595/4128690/GridPPDurhamXcacheVP_Mar2022.pdf
- Data from Oxford showing that reactive XCache not the most efficient
https://indico.cern.ch/event/1128343/contributions/4787170/attachments/2412595/4128683/20220323_XcacheMaterial_draft2.pdf
- Working to get S3 endpoints working

Monitoring of the Data Lake

- Create a series of metrics which will measure: Overall time for [data analysis] tasks to complete, storage used, bandwidth used, CPU efficiency of tasks and operational burden (sites and experiment). For comparison with current model.
- Rucio monitors storage used for each RSE
- Currently using FTS monitoring for the bandwidth and transfer efficiency
- Moving to K8S plan to incorporate monitoring across Rucio as well



Future Developments

- Move Rucio to a K8S deployment
- Improve monitoring for Rucio
- Configure of additional sites and technologies:
 - Cloud storage
 - cache for onsite analysis
 - XCache
- Implementation of QoS
 - Some work has been done within Rucio to support this
 - Still to be looked at and completed

Summary & Questions

- Data-lake setup with a number of sites
 - Working to expand the data-lake
 - Looking forward to integration with workflow management
 - How best to do this? Which VO/s?
 - XCaches and how to use them to offer the best performance?
- Monitoring for Rucio
 - Using Rucio & FTS to monitor storage and transfers
 - Expand Rucio monitoring with K8S deployment
 - Should overall monitoring of data and workflow be unified?
- QoS for data placement
 - QoS still a work in progress?
 - How would this work with how Endpoints are set up now?
 - What may need to change?