

Rucio Meeting notes 2023-05-01

Date

01 May 2023

Zoom

<https://stanford.zoom.us/j/91889629258?pwd=KzNleVdmSnA1dkN6VkRVUTNtMHBPZz09>

Attendees


- [Richard Dubois](#)
- [Brandon White](#)
- [Kian-Tat Lim](#)
- [Brian Yanny](#)
- [Yuyi Guo](#)
- [Steve Pietrowicz](#)
- [Wei Yang](#) (chair)
- [George Beckett](#)
- [Greg Daues](#)
- [Andy Hanushevsky](#)
- [Matt Doidge](#)
- [Wen Guan](#)
- [Peter Love](#)
- [Timothy Noble](#)
- [Lionel Schwarz](#)
- [Peter Clark](#)
- [Michelle Gower](#)
- [Fabio Hernandez](#)

Apologies

Agenda — Data Replication

- Status of replication exercises among the 3 facilities [[Yuyi Guo](#)[Brandon White](#)]
 - update on the status of the FTS instances (Rubin's and UKDF's)
 - update on the replication exercises using Rubin's Rucio and UKDF's FTS. Is the issue of simultaneous transfers solved?
 - what is preventing us to start doing regular replication of the equivalent of 1 night's worth of raw data, first using FTS only, then Rucio?
- Status of Rucio & FTS monitoring [[Timothy Noble](#) [George Beckett](#)]
 - update on the improvements to the Rucio and FTS monitoring platform
- Status of butler & Rucio integration [[Steve Pietrowicz](#)]
 - deployment progress at the DFs of the consumers of Rucio events. Relevant JIRA issues:

■  [DM-38655](#) - Jira project doesn't exist or you don't have permission to view it.

■  [DM-38656](#) - Jira project doesn't exist or you don't have permission to view it.

■  [DM-38657](#) - Jira project doesn't exist or you don't have permission to view it.

- since the Kafka topics are to be named after the names of the RSEs, shouldn't we start standardizing our naming conventions to the official naming used by the project (i.e. USDF, UKDF, FrDF)?
- Status of replication of STS data from USDF to FrDF [[Wei Yang](#)[Kian-Tat Lim](#) [Fabio Hernandez](#)]

Data Replication JIRA Issues

- JIRA tickets with tag "data-replication"

▪ **Key** **Summary** **T** **Created** **Updated** **Due** **Assignee** **Reporter** **P** **Status** **Resolution**


No issues found

Notes

Data Replication (Fabio)

- Data Replication Experiments
 - Dan Speck helped Brandon to deploy FTS database. Hope to be ready for testing within two weeks (bootstrapping and testing).
 - Brandon has been working on data-replication plan.
 - Steve P noted worth Brandon talking to him when testing the database, as need configuration for ActiveMQ and Kafka messaging (should be done automatically).
 - Yuyi noted currently using FTS instance at RAL, though have experienced issues with simultaneous transfers. Yuyi and Brandon have confirmed how to set files per job and have successfully tested 100 files per job on SLAC FTS. In transfer test with RAL FTS, started with 100 files per job but dropped back to 1 file per job after a while. Difficult to debug issue at RAL FTS.
 - Rucio cannot do bulk transmission of files from different sources. Required an appropriate distance schema to ensure transfers resolved down to one source (not two).
 - KT asked if experience suggested would be issues with Rucio getting confused when needing to transfer data back to USDF, for example.
 - Yuyi noted need to manipulate data to ensure transfer happens from SLAC to Europe (with existing dataset, Rucio chooses closer instance of file in Europe)
 - Greg proposed to set up subscriptions and rules with explicit end-point definition : does not see a situation when would not know where data from moving from or to.
 - RAL FTS has a library problem, preventing current use.
 - Matt confirmed seeing other issues with intersite transfers, since RAL updated FTS during week beginning 24th April.
 - Rose Cooper (RAL) has joined Slack channel to help resolve issue
 - Fabio proposes that RAL FTS status is logged in Jira ticket (Greg offered to transcribe notes from Slack to Jira)
 - Richard noted expectation to send Lite catalogues to around 12 IDACs, and would prefer not to send 12 copies from SLAC.
 - Fabio confirmed not yet ready to undertake large-scale transfer campaigns.
- Monitoring
 - Waiting on FTS at SLAC to set up further monitoring in Grafana.
- Butler-Rucio Integration
 - Steve P to finalise testing plus containerise for ingestion side, following suggested approach from K-T
 - For Rucio daemon (Hermes) side is potentially more complicated. Hermes container configuration looks complicated
 - Potentially deployment significantly simpler if moved to newer version of Rucio, plus to make Hermes daemon better isolated
 - Brandon reluctant to move away from LTS version, which has been selected.
 - Next LTS is due in July (Version 32)
 - Steve P will share some notes on potential disadvantage of current LTS version.
- Service deployment
 - Jira tickets created detailing what services need to be deployed at each site.
 - Fabio suggested to standardise on naming for site-based Rucio services, to avoid strange names persisting for longer term.
 - Brandon believed would be trivial operation to change naming of Rucio services
 - K-T noted names recorded in various directories, so worthwhile to check their suitability in short term.
- STS data replication from USDF to FrDF
 - KT confirms registration of data is working again (after problem resulting from power outage)
 - Still to test automatic subscription (or to setup rule to register file when transferred)
 - Brandon suggested KT talk to Yuyi about this
 - Also need to setup registration of data from camera with Rucio

Multi-site PanDA Support and Operation (Wei)

- Brian noted update on progress for large-scale test of multi-site production
 - MS to demonstrate scaling capability of multi-site processing by end of July.
 - Propose to use PDR2 (17k exposures from HSC) – 50 TB of raw inputs to progress through Step 1 processing at USDF to create calibrated images to be transferred to UKDF and FrDF (100 TB split between two sites) Step 3 (Coadd processing) run at data sites to be returned to USDF (around 50 TB of outputs)
 - Progressing beyond Step 3 would be a stretch goal: Step 3 is a good test of memory requirements.
 - Wei noted several limitation
 - Manual ingestion of data into Butler
 - Need to resolve large-sale Rucio transfers
 - Intermittent issues running on USDF (compared to French DF) possibly due to NAT
 - Need to refine memory requirements (potentially setting up multiple PanDA queues to address requirements)
 - Work with Wen to set up queues (having confirmed schema for queues and memory requirements)
 - Some issues being investigated with memory requirements of steps in testing.
 - Need to work out how to resubmit failed job with increased memory, if one fails
- PanDA (USDF installation)
 - Wen noted currently working on development environment
 - PanDA server is now running independently of CERN IAM and CRIC  server
 - Scaling test, with 3,000 cores, went okay (failure rate acceptable)
 - When jump to 6,000 cores see too many failures, possibly related to upload of logs to storage (bottleneck in Squid Proxy)

- Need Kubernetes allocation to progress deployment of production PanDA instance

Date of next meeting

Monday May 15th, 8am PT

Discussions on Slack [#dm-rucio-testing](#) in interim

Action items

