

# Rubin Observations

A quick blast of things we're doing, or have achieved, noticed or had to hack around:

- Rubin are currently using the RAL FTS with success
  - The test instance at SLAC was not up to spec.
- The equivalent of a full night's data has been replicated to Lancs (from IN2P3)
  - Lots of small files, it caused more load than you'd expect from a few TB.
  - Was not done using rucio.
- Peter Love has been running test pipeline jobs on this data using panda.
  - Some base level debugging - jobs not requesting enough RAM/time.
  - Interaction with butler credentials means rubin users restricted to one user, rather than a pool.
- Lancaster preparing to set up messenger containers, but waiting on feedback.
- Some recent leaps forward in the panda pipeline at SLAC (according to messages on slack).
- There is a milestone for a 3-site data processing run by the end of July
  - Looks good to achieve this, although there might be more manual interventions than with the finished product.