



ASTRON

Netherlands Institute for Radio Astronomy

LOFAR use case in ESCAPE


Experience and future directions
Yan Grange, Vishambhar Nath Pandey, Klaas Kliffen


ASTRON

Netherlands Institute for Radio Astronomy




LOFAR in one slide

 The image part with relationship ID rld5 was not found in the file.

 The image part with relationship ID rld5 was not found in the file.

From the first stars to lightning on earth,
and anything in between



ESCAPEuse cases for DAC21:


Context

- The LOFAR use cases are based on the datalake as the Long-Term Archive (LTA) of the telescope
 - But also maybe as a way to manage several archives of different telescopes
- Current LTA is distributed over three locations, most data management happens manually
 - Replication is currently not really an issue (we can live with some data loss) but it may be a very interesting concept for popular data or end products.
- Data to be embargoed, only accessible for the requestor for 1 year, after which it becomes fully public (=>general implications on AAI)
- LTA data sizes: 7PB/yr originally. Now adding more processing and compression: 5PB/yr. But we are upgrading our telescope, so it is likely to go back up again. Also new modes get added which could incidentally create larger data sets.




ESCAPE use cases for DAC21:

Ingestion and replication



- Goal is to mimic a day of ingest of LOFAR data.
 - 2 observations: 1 target, 1 calibrator.
 - Typically 15TB of data in 400 files (~60G per file for target, 2GB per file cal).
 - Files are actually tarred directories
- Using custom tooling to “manually” upload the data (to a non-deterministic RSE), then register to Rucio for QoS life cycle
 - Data is primarily stored on tape (I call that “slow” or “cheap” QoS)
 - Data is kept on disk for a week or so (i call this “direct accessible” QoS)



LOFAR


ASTRON

Netherlands Institute for Radio Astronomy

ESCAPE use cases for DAC21:

Data Processing

- Goal is to process data that is on the data lake in a Jupyter lab notebook to an image
- Assume data has been uploaded to the data lake before (ie was uploaded before DAC21)
 - Create data hierarchy: 2 targets, 1 calibrator. All files for a single observation in a Dataset.
 - Each target dataset in a collection with the calibrator dataset
- Query data in ESAP (science analysis platform), use ESCAPE-developed DataLake as a Service notebook to create an image.
 - Multiwavelength: combine the image made from data in the datalake with images from other wavelength registered in the Virtual Observatory (ie global astronomy archive of data)
- Upload processed image back to the datalake



ESCAPE use cases for DAC21:

Legacy archive

- Goal is to simulate the case where we would go for Rucio, and want to add our current archive as a “legacy data source”.
- Add one current LTA location as an RSE in the ESCAPE Rucio
 - AAI may be a challenge here. Current LTA has a different VO from ESCAPE, but this may be a realistic problem anyway 😊



LOFAR

ASTRON

Netherlands Institute for Radio Astronomy

Experiences:

Disclaimer: no operational experience, few locations, etc. Real life processing of a few (23) freq. subbands of data.

Most experience stems from the ESCAPE Final Dress Rehearsal, and some experimenting for DAC21.

1. Getting data from the LTA into the datalake
 - a. We've transferred some data from an external location (the current LTA) to the ESCAPE datalake. Moving between gridftp in the VO "lofar:/lofar/user" to a gridftp end point in "escape:/escape/lofar" was a bit hard (not really a Rucio issue).
 - b. Very basic experience with AAI otherwise.



LOFAR

ASTRON

Netherlands Institute for Radio Astronomy

Experiences:

Disclaimer: no operational experience, few locations, etc. Real life processing of a few (23) freq. subbands of data.

2. Upload/download

- a) Single files can be quite big. We have experience that some tooling may time out when the computation of the checksum takes longer than the network timeout. Noticed it is a bit hard to debug because of all the parts in the stack (Rucio, calling gfal calling a specific copy tool).
- b) Downloads of hierarchical data all end up in same directory, ideally would like to keep the hierarchy in the download
- c) DLaaS is a very nice development.

3. OIDC rocks (or: I don't really like X509)!

- a) LTA has historically been on X509-backed storage and that has been quite an issue for our users (since we happen to only have dCache backed locations tokens are nowadays possible).
- b) I guess this is especially true for a public archive



LOFAR

ASTRON

Netherlands Institute for Radio Astronomy

Dreams and hopes (AAI)

Rucio seems to have lots of potential for our use cases

- AAI granularity
 - e.g.: public user, group of owner, data owner (PI), observatory staff, admin
 - Also granularity in who can grant access to data:
 - The PI can grant (either directly or by proxy) access to data, observatory stadd/admin to PI.
- AAI based on metadata
 - ie: if file is part of a certain project, AAI of that project is followed.

Dreams and hopes (Data and transfers)

- Debugging where transfers time out/fail, and how to set timeouts or other transfer settings
- MS format on tape means having to tar it. This causes quite some overhead. Having the tar step "pipelined in" the transfer would be very nice, as would be (un)tarring at the remote site (tarred as a QoS level?)
- Data hierarchy is very common in astronomy. Would be nice to keep the hierarchy in the datalake in the locally downloaded copy (e.g. after download I get *MYSCOPE:containerX/datasetA/fileU.tar*).
 - Also it is not so clear from our perspective what else than opacity a deterministic RSE would add.



LOFAR

ASTRON

Netherlands Institute for Radio Astronomy

Dreams and hopes (File system)

- Direct remote access to files as if they were mounted as POSIX, downloading only parts of files (e.g. headers, or a cone in an image)
- Integration between iRods (as an RSE?) and Rucio. Maybe even communication of rules between the two?



LOFAR

ASTRON

Netherlands Institute for Radio Astronomy

Dreams and hopes (Metadata)

- Custom metadata that is the same in the user-facing queryable database, the Virtual Observatory and the Datalake. Direct connection between Rucio and observation database for shared metadata.
- Advanced metadata rules, including simple math ($a^{**2}+b^{**2} < X$ as simple cone example). Or setting where dat will go based on projectID.



LOFAR

ASTRON

Netherlands Institute for Radio Astronomy

Dreams and hopes (up- and download)

- Integration of a local filesystem as RSE, even if no root access.
 - Typical use case: automation of ingest. Telescope writes to disk, notifies Rucio and replication rule (target location based on metadata) is triggered, and the disk space is vacated when done.
- Multihop download (academic networks can be... complex)
 - Download via proxies,
- The two above points combined (e.g. an RSE that is behind a restrictive firewall and needs some magic to get data in or out).
- Multi-step download (e.g. "Stage data from tape and when that is done, download it to local disk"). Is there a link with file system messages here?



LOFAR

ASTRON

Netherlands Institute for Radio Astronomy

Dreams and hopes (other)

- Monitoring of the data lake, including the network, so that problems can be spotted easily
- Integration with science data repositories (DOI de-referencable data endpoint, pre-staging if need be)



LOFAR

ASTRON

Netherlands Institute for Radio Astronomy