Rucio and the International Gravitational Wave Network

James Alexander Clark, LIGO

james.clark@ligo.org, 7th Rucio Community Workshop, LIGO-G2402116

International Gravitational Wave Network

Observing status

Data taken in observing runs of months' - years' stable instrumental (and software) configuration; Current O4 (start May 24, 2023) status (see https://gracedb.ligo.org/superevents/public/O4/):

- O4 Significant Detection Candidates: 140 (156 Total - 16 Retracted)
- O4 Low Significance Detection Candidates: 2459 (Total)

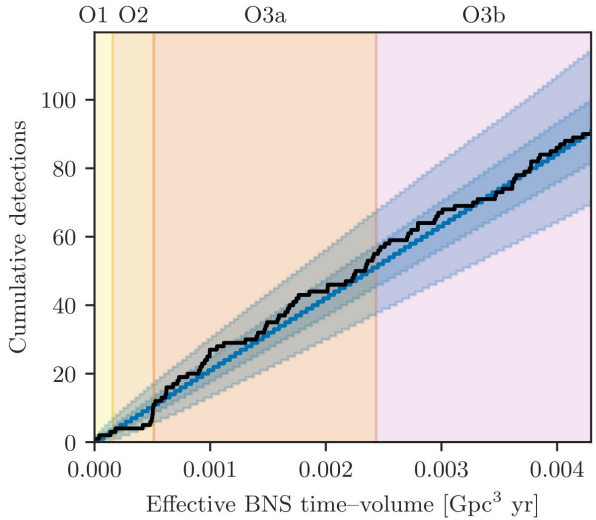

90 candidates observed O1-O3
arXiv:2111.03606

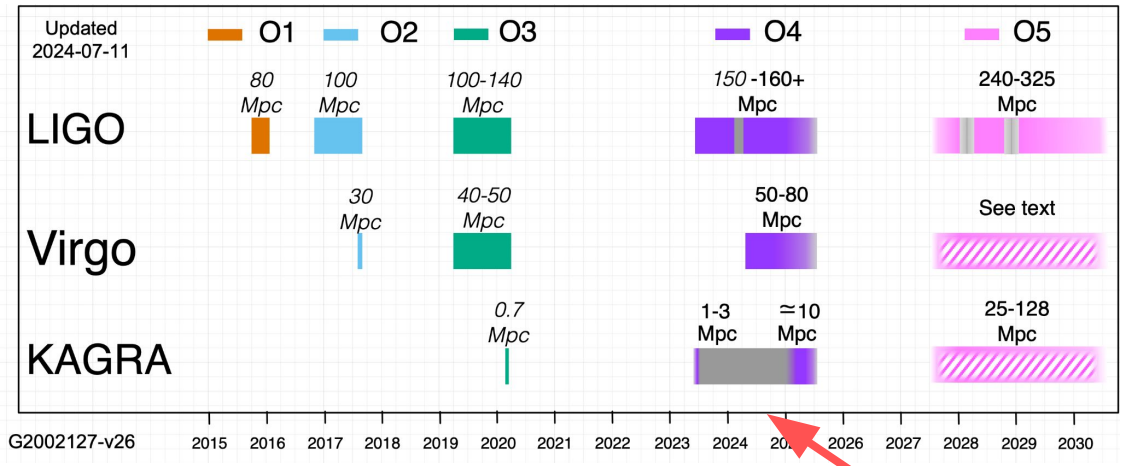

https://observing.docs.ligo.org/plan/

We are here

Data products & management

Instrumental data from O(500,000) channels in LIGO/Virgo instruments recorded into “frame files”:

- Low-latency data streams: ~real-time searches & public alerts (finite ~weeks lifetime)
- **High-latency bulk frame files**: instrument data aggregated into various frame file types at different levels of reduction → long-term custodial storage & offline / archival searches, computationally expensive parameter estimation etc

Rucio used in LIGO to manage high-latency streams / custodial storage.

Bulk data → dHTC resources via OSDF

~10 different LIGO datasets with file cadence between 64-4096s (~few GB per file)

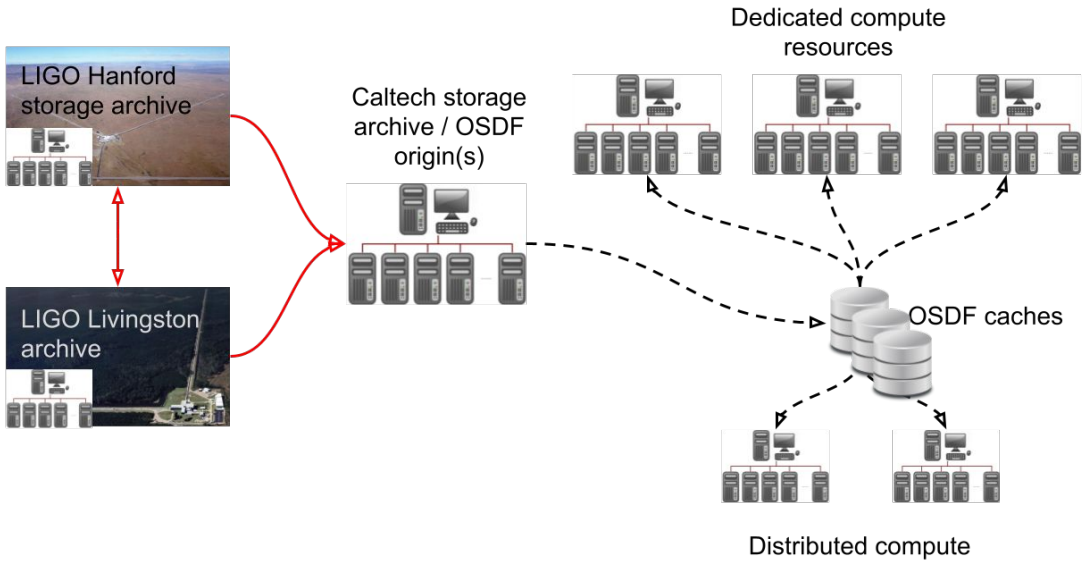

Data products & management

O4 archival frame file data sets (since May 24 2023)

Frame type	Total Size	Count	Typical size / file	Typical duration / file (s)
Raw	2.4P	1.3M	2G	64
Strain	26.9T	21.0K	1.6G	4096
Strain ("analysis ready")	17.5T	14.3K	1.6G	4096
Minute-trend	11.5T	23.7K	500M	3600
Second-trend	141.7T	142.4K	1.1G	600

Still managed by legacy LIGO Data Replicator system

- Transition for highest priority set pending backup operator training (in progress)
- Entire deployment, admin, development & operation effort to date → <100% of my time.

GravCat

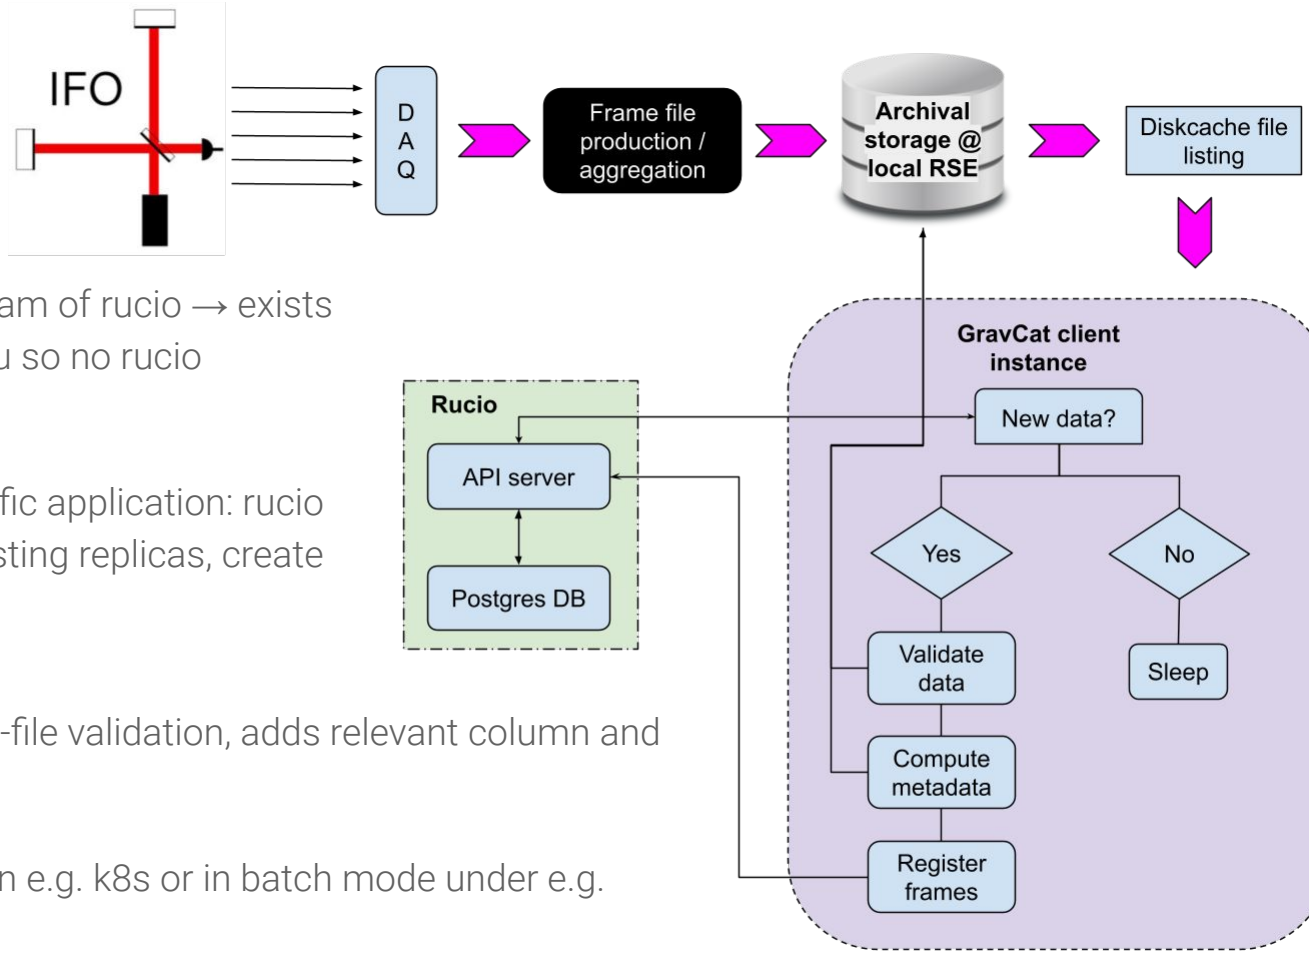

Data is produced upstream of rucio → exists in archival storage in-situ so no rucio uploads.

“GravCat”: domain-specific application: rucio client API to register existing replicas, create datasets etc

Performs internal frame-file validation, adds relevant column and JSON metadata.

Runs in daemon mode in e.g. k8s or in batch mode under e.g. HTCondor.

LIGO Rucio deployment architecture

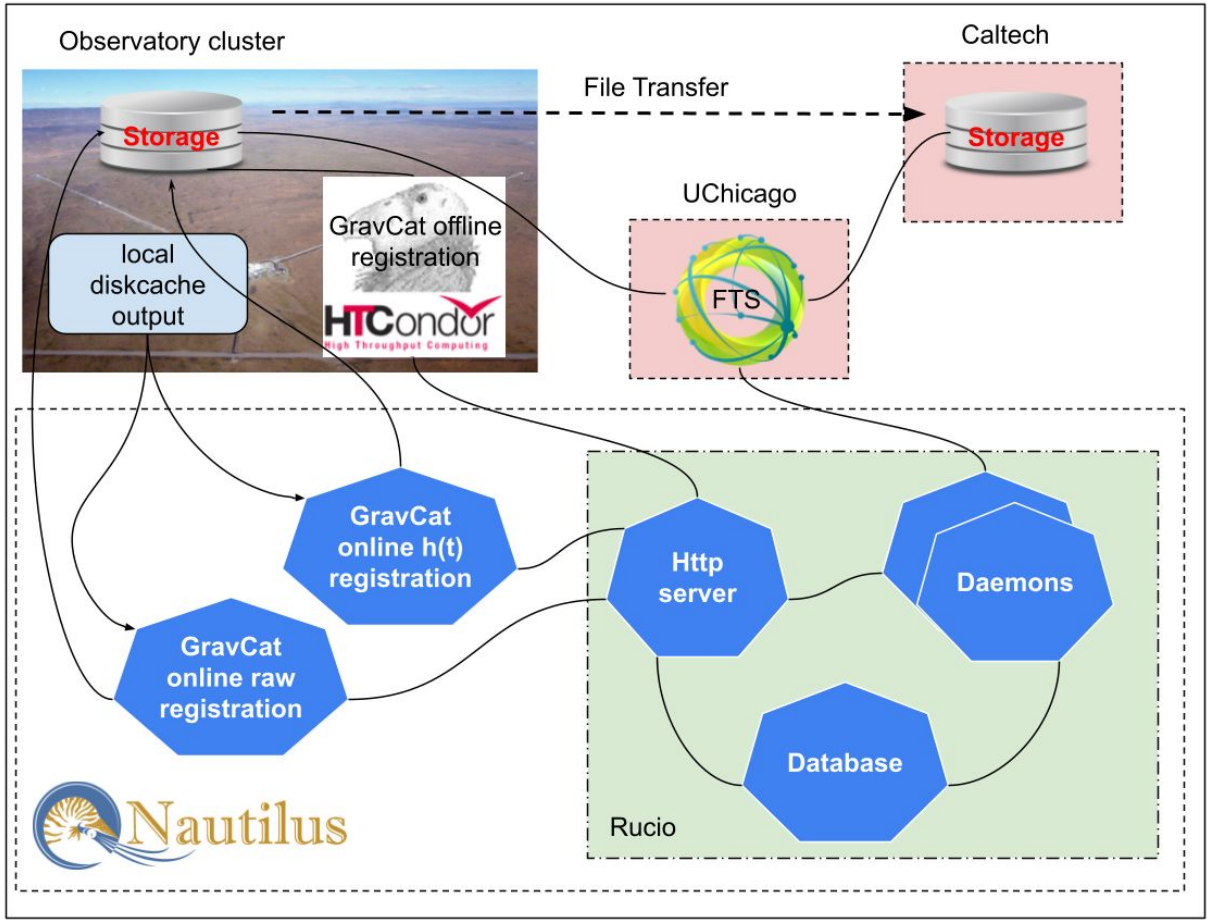

Some conventions: scopes, LFN2PFN, ...

Desired POSIX paths for IGWN frame files generally determined by: a) observing run & a basename convention.

E.g. LFN: `H-H1_HOFT_C00-1411305472-4096.gwf` from observing run **O4** has desired PFN:

`<local RSE prefix>/O4/hoft_C00/H1/H-H1_HOFT_C00-141/H-H1_HOFT_C00-1411305472-4096.gwf`

So: deterministic LFN2PFN in custom policy package & YAML lookup table for data type translations

```
- frame_type: "HOFT_C00_AR"
  generic_type: "hoft_C00_AR"
  gps_truncation: 3
```

Encode observing run in the scope like `<account>.<data-product>.<observing-run>`, e.g.: `frames.L1-hoft.O4`

Can support admins' preferences for file distribution compatibility evolution of PFNs conventions over years of observing runs.

Subscriptions

Use Rucio *datasets* to represent each parent directory of frame files, e.g. `H-H1_HOFT_C00-141` on previous slide:

`frames.H1-hoft.O4:H-H1_HOFT_C00-141`

Define subscriptions based on GravCat-imposed metadata and scopes:

```
$  rucio-admin subscription list O4-CIT-hoft-ar
frames: O4-CIT-hoft-ar UPDATED
  priority: 3
  filter: {"stream_name": ["L1_HOFT_C00_AR", "H1_HOFT_C00_AR"], "did_type": ["DATASET"], "scope": ["frames.L1-hoft.O4", "frames.H1-hoft.O4"]}
  rules: [{"copies": 1, "rse_expression": "LIGO-CIT", "activity": "User Subscriptions"}]
```

Resulting rules:

```
$  rucio list-rules --subscription frames O4-CIT-hoft-ar
ID                                ACCOUNT    SCOPE:NAME                                    STATE[OK/REPL/STUCK]    RSE_EXPRESSION    COPIES    EXPIRES (UTC)    CREATED (UTC)
--------------------------------  ---------  --------------------------------------------  ----------------------  ----------------  --------  ---------------  -------------------
b7c4b76838b84741adfc5c8ffe2f4388  frames     frames.H1-hoft.O4:H-H1_HOFT_C00_AR-137  OK[2115/0/0]            LIGO-CIT          1                          2023-09-11 19:06:52
e01cb4dfcdb8408da4e002465b601ae6  frames     frames.H1-hoft.O4:H-H1_HOFT_C00_AR-140  OK[1177/0/0]            LIGO-CIT          1                          2024-05-18 01:34:38
cb97aa2c8b9d4cafa540f17cd4a8528c  frames     frames.H1-hoft.O4:H-H1_HOFT_C00_AR-136  OK[152/0/0]             LIGO-CIT          1                          2023-09-11 19:06:51
191d4ec07656482784201bfccc06b7cb  frames     frames.H1-hoft.O4:H-H1_HOFT_C00_AR-141  OK[260/0/0]             LIGO-CIT          1                          2024-09-10 13:01:37
d4407b2c852043b4b3661039466578ae  frames     frames.L1-hoft.O4:L-L1_HOFT_C00_AR-141  OK[227/0/0]             LIGO-CIT          1                          2024-09-10 13:21:39
670d2778ce6347e88691b7722de17d8a  frames     frames.L1-hoft.O4:L-L1_HOFT_C00_AR-140  OK[2240/0/0]            LIGO-CIT          1                          2024-05-17 18:52:15
012a25b83c4d43c0b70d1cca55a08aca  frames     frames.L1-hoft.O4:L-L1_HOFT_C00_AR-138  OK[1594/0/0]            LIGO-CIT          1                          2023-09-29 08:20:16
941d34a2f4a441e0933fd4b6787f606e  frames     frames.H1-hoft.O4:H-H1_HOFT_C00_AR-138  OK[1919/0/0]            LIGO-CIT          1                          2023-09-29 08:16:16
e0ded56e84db421bb660c8ef86ded65c  frames     frames.L1-hoft.O4:L-L1_HOFT_C00_AR-139  OK[1027/0/0]            LIGO-CIT          1                          2024-03-21 02:23:51
8b59c92307684b4c93ffc77be1b88bf3  frames     frames.H1-hoft.O4:H-H1_HOFT_C00_AR-139  OK[940/0/0]             LIGO-CIT          1                          2024-03-21 02:03:49
```

Monitoring

Use native Rucio events → ELK stack to monitor general health, quantify activity.

Primary health indicator: legacy (20+ year) “lag” monitoring.

Cronjob compares time-stamps of LFNs with current time at destination.

Visual inspection for obvious problems.

Also use Icinga lag monitoring for alerts.

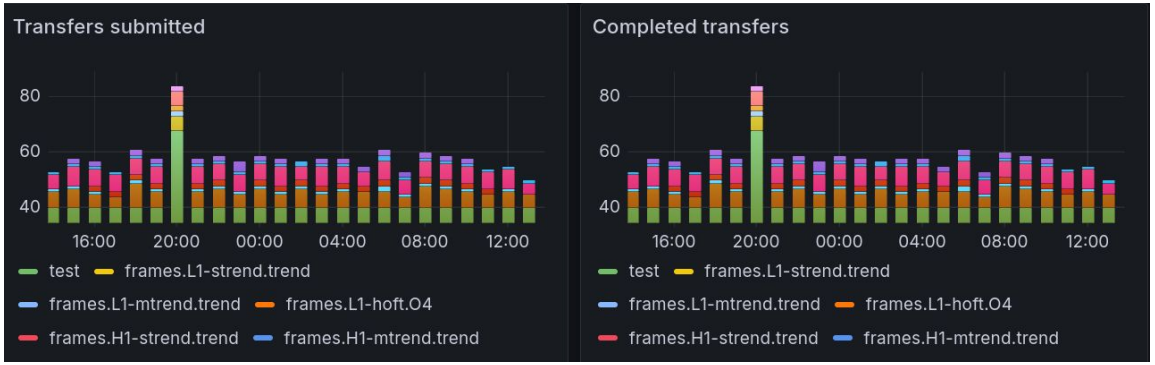

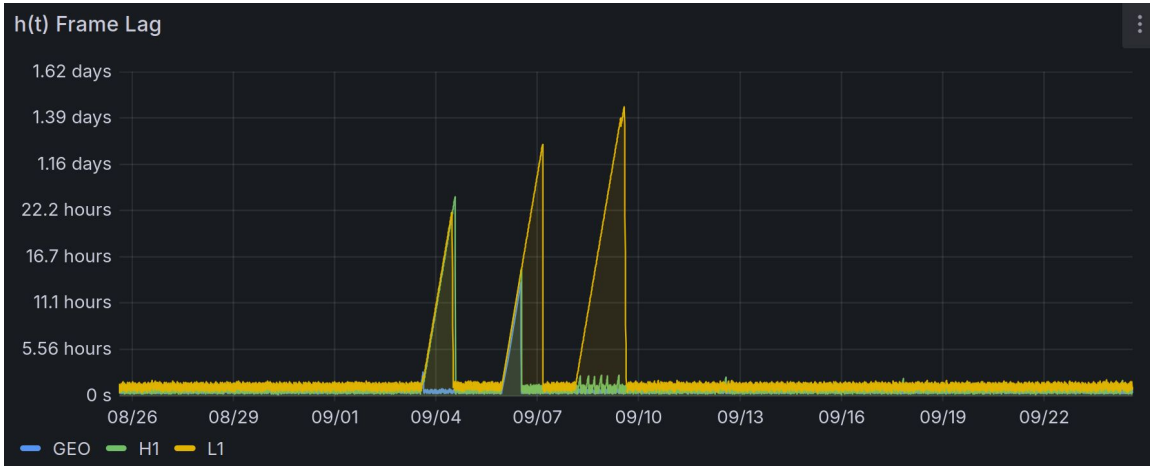

Virgo Rucio Setup

- Virgo moved to Rucio in 2020 to replace legacy replication software
- Data flow:
  - *Production:* h(t) + env condition data transferred via Rucio from EGO to Origin Stashcache in Louvain, and then distributed through CVMFS/OSDF
  - *Under Test*: raw transfer from EGO to INFN Turin

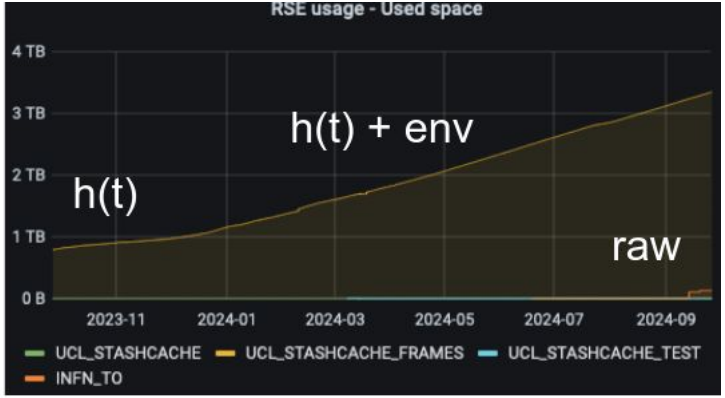

- production and test Rucio servers+DB deployed at INFN Turin
- custom client application running at EGO to upload data to Rucio

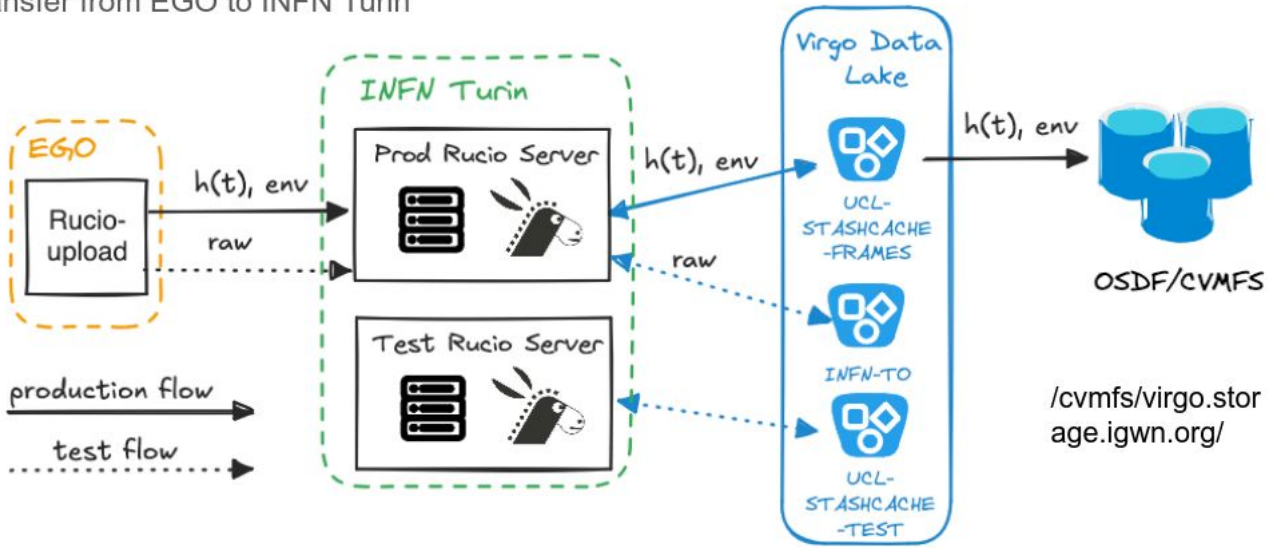

Feedback

No out-of-the-box data discovery tools and non-trivial in-situ file registration

- likely very domain specific so probably not a reasonable ask (?)

Desirable to extend metadata support

- Can repurpose existing attributes but subscriptions and native metadata would be ideal.
- e.g. metadata matches based on time/frequency/… inequalities.

Rules sometimes remain stuck after an outage

- Sometimes need to boost a rule to un-stick it after end-point or FTS outages

FTS doesn't use temporary filenames 😡

- if connectivity gets bad enough during a transfer FTS fails to clean up a failed transfer.
- Incomplete files can then be exposed to users / archive processes.
- Unclear what *Rucio* could reasonably do about this.

Finally:

- Fantastic development/user community!
- Excited to learn more, understand scope to extend to other data management tasks.