Rucio & dCache in the SDC@KIS

Hosting inhomogeneous solar data on multiple sites

5th Rucio WS, Lancaster

Leibniz-Institut für Sonnenphysik (KIS)

P.Caligari, M. Günter, A. Hochmuth

J. Beck, S. Berdyugina, A. Bührer, V. Gangadharan, N. Bello Gonzalez, A. Gorobets, L. Hohl, P. Kehusmaa, M. Knobloch, S. Müller, S. Patel, T. Yakobchuk

Background

Solar Observational Data

Versatility: bone and bane of ground-based solar physics.

- Highly inhomogeneous data
- Existing data is not standardized
- Tailored post-processing needed

⇒ 2D/3D maps of physical parameters ($\vec{B}$, $\vec{v}$, P, T, …) as a function of time

- 3 – 70 TB/d

Schlichenmaier et al. (2016)

Naming & data injection

Leibniz-Institut für Sonnenphysik (KIS)

- Existing data files are named VERY arbitrarily and stored in a custom directory tree
  - `BBI: bbi_20170519-081224_000_WL656_reco_bb.fits`
  - `LARS: lars_l20_20161018-114421_clv6302_mu06_w.ns.chvtt.sum.R100000.FeI6302.5.bisect.fits`
  - `GRIS: 02jun21.007-02`
  - Lots of instrument specific fields
  - Requires in-depth knowledge
- Rucio naming
  - Roughly based on SOLARNET recommendations (https://arxiv.org/abs/2011.12139)
  - Keep for L0-data (raw & calibration files), change for science-ready data (L1, L2, L3) to

```
campaignID|instrument:
instrument[-mode/subtype][_L#]_YYYYMMDD[THHMMSS[FFF]][[_misc]_contenttype]][_V#].filetype
```

- Multi-instrument importer framework in Python (kharon)
  - *Observation*: the logical viewing unit (instrument dependent!) ≙ Rucio containers

Metadata

Grouping: *observations*

- kept in MongoDB
- Corresponds to Rucio *container*
- Header: IVOA
- Used for VOs, e.g.: https://solarnet.oma.be

Science-ready data

- kept in Rucio (file)
- Versioned!
- FITS-files[2)] (contain SOLARNET conform header[1)])
- Subject to embargoes
- Header: extended FITS header
- Eventually associated to DOIs

Volatile data
(previews, log-files, calibration files, etc.)

- Kept in MongoDB / GridFS
- **No** science-ready data
- Used on: https://archive.sdc.leibniz-kis.de
- One (latest) version, only!

[1)] https://arxiv.org/abs/2011.12139.
[2)] https://fits.gsfc.nasa.gov/fits_documentation.html

Metadata

Grouping: *observations*

- kept in MongoDB
- Corresponds to Rucio *container*
- Header: IVOA
- Used for VOs, e.g.: https://solarnet.oma.be

Rucio DID		MongoDB		File
Metadata	=	Header	=	FITS-Header

Science-ready data

- kept in Rucio (file)
- Versioned!
- FITS-files[2)] (contain SOLARNET conform header[1)])
- Subject to embargoes
- Header: extended FITS header
- Eventually associated to DOIs

Volatile data
(previews, log-files, calibration files, etc.)

- Kept in MongoDB / GridFS
- **No** science-ready data
- Used on: https://archive.sdc.leibniz-kis.de
- One (latest) version, only!

[1)] https://arxiv.org/abs/2011.12139.
[2)] https://fits.gsfc.nasa.gov/fits_documentation.html

Metadata

Leibniz-Institut für Sonnenphysik (KIS)

Grouping: *observations*

- kept in MongoDB
- Corresponds to Rucio *container*
- Header: IVOA
- Used for VOs, e.g.: https://solarnet.oma.be

gris_observations

- **_id: ObjectID(name_uid)**
- metadata: dict
  - header: dict (IVOA)
  - tarsize_l1data: int
  - data_quality: int
  - campaign_id: str
  - embargo_until: datetime
- links: dict
  - type1: list of DBRef
  - ...
  - typeN: list of DBRef
- lastModified: datetime

gris_data

- **_id: str (DID)**
- filename: str
- length: int
- metadata: dict
  - date: datetime
  - type: str
  - header: dict (opt, SOLARNET)
  - md5: str
  - mimeType: str
- uploadDate: datetime

gris_auxdata (GridFS)

- **_id: ObjectID**
- filename: str
- length: int
- chunkSize: int
- metadata: dict
  - date: datetime
  - type: str
  - header: dict (opt)
  - md5: str
  - mimeType: str
- uploadDate: datetime

Science-ready data

- kept in Rucio (file)
- Versioned!
- FITS-files[2)] (contain SOLARNET conform header[1)])
- Subject to embargoes
- Header: extended FITS header
- Eventually associated to DOIs

Volatile data
(previews, log-files, calibration files, etc.)

- Kept in MongoDB / GridFS
- **No** science-ready data
- Used on: https://archive.sdc.leibniz-kis.de
- One (latest) version, only!

Rucio DID → Rucio MD plugin → MongoDB Header = File FITS-Header

[1)] https://arxiv.org/abs/2011.12139.
[2)] https://fits.gsfc.nasa.gov/fits_documentation.html

Embargoes

- X509 based (certificates only valid for 1 year!); no VOs
- Several asynchronous daemons per RSE involved
- Is this concept still feasible after switching to token-based AAI (IAM)?
  - Should we wait?
- File-based, not scope-based! Would require:

```
rucio list-dids*:instrument*
```

SDC Sizing (11.2022)

Tier1 @KIS

- All data (L0, L1, L2, L3)
- User caches
- 2 copies per file (dCache)
- Lifetime: years

Tier 2 @KIS

- All data
- **No** user caches
- 1 copy per file
- Lifetime: > 11 years

Tier1 @OT

- Raw data (L0)
- **No** user caches
- 1 copy per file
- Lifetime: weeks/months

Next Steps (2023)

Leibniz-Institut für Sonnenphysik (KIS)

- Integration with telescope interfaces at OT (associate data to campaigns)
- Fully automating data transport from OT to KIS@SDC
  - Establish a (semi-) automatic data quality selection
  - Automated data calibration (L0 -> L1) in SDC@KIS
- Integrate
  - Astronomical VO and SOLARNET VO
  - Long term: join ESCAPE data lake (use IAM) & join Punch4NFDI data lake
- Inject first non-OT data

SOLARNET ESCAPE PUNCH 4NFDI

- Test and establish upgrade and disaster recovery procedures for dCache and Rucio
- Rucio & dCache monitoring

- Still struggling to put all the pieces (FTS, dCache, daemons, k8s, …) together

[1] PUNCH4NFDI = consortium of article, astro-, astroparticle, hadron and nuclear physics of the German National Research Data Infrastructure

Open Questions

Leibniz-Institut für Sonnenphysik (KIS)

- JSON-blobs in PostgreSQL seem slow?
  - Way to speed up other than using an external database for custom MD?
  - Rucio & MongoDB have to be kept in sync!
- Unstable internet connection to Canary Islands
  - Rucio is located in Freiburg, but data is injected on Tenerife
    ⇒ 2-step injection process (copy to RSE, then register with Rucio)
  - Distributed Rucio server with multiple instances (e.g. using PostgreSQL Multi-master Replication)?
- ESCAPE <-> Punch4NFDI
  - SDC is involved in both (in ESCAPE via EST)
  - E.g. AAIs are compatible?
- More flexible fine-grained (Posix-like) permissions **per file**?
- Slack -> paid version? (- > Mattermost, RocketChat, etc…)
- We are still stumble upon networking issues in k8s cluster and running the different components.