



ESCAPE

European Science Cluster of Astronomy &
Particle physics ESFRI research Infrastructures

The ESCAPE Project: Data Lake and Science Platform

Yan Grange
grange@astron.nl

Klaas Kliffen
kliffen@astron.nl

John Swinbank
swinbank@astron.nl



Introductions & ESCAPE Overview

John Swinbank



Overview

- Introductions: who are we?
- Context: the ESCAPE project, its environment, its acronyms.
- Other ESCAPE deliverables: software repository, VO, citizen science
- Deep dive: DIOS — the “Data Lake”
- Deep dive: ESAP — the “Science Platform”



ESCAPE Goals & Fast Facts

- **Develop common “e-infrastructure” solutions that benefit a wide range of particle physics & astronomy research facilities.**
- ~30 different partners across 8 countries.
- Operating in partnership with four other “cluster” projects covering environmental research (ENVRI), neutron & photon science (PaNOSC), life sciences (EOSC-Life), and social sciences and humanities (SSHOC).
- Project runs from 2019 until 2023.



ESCAPE Consortium



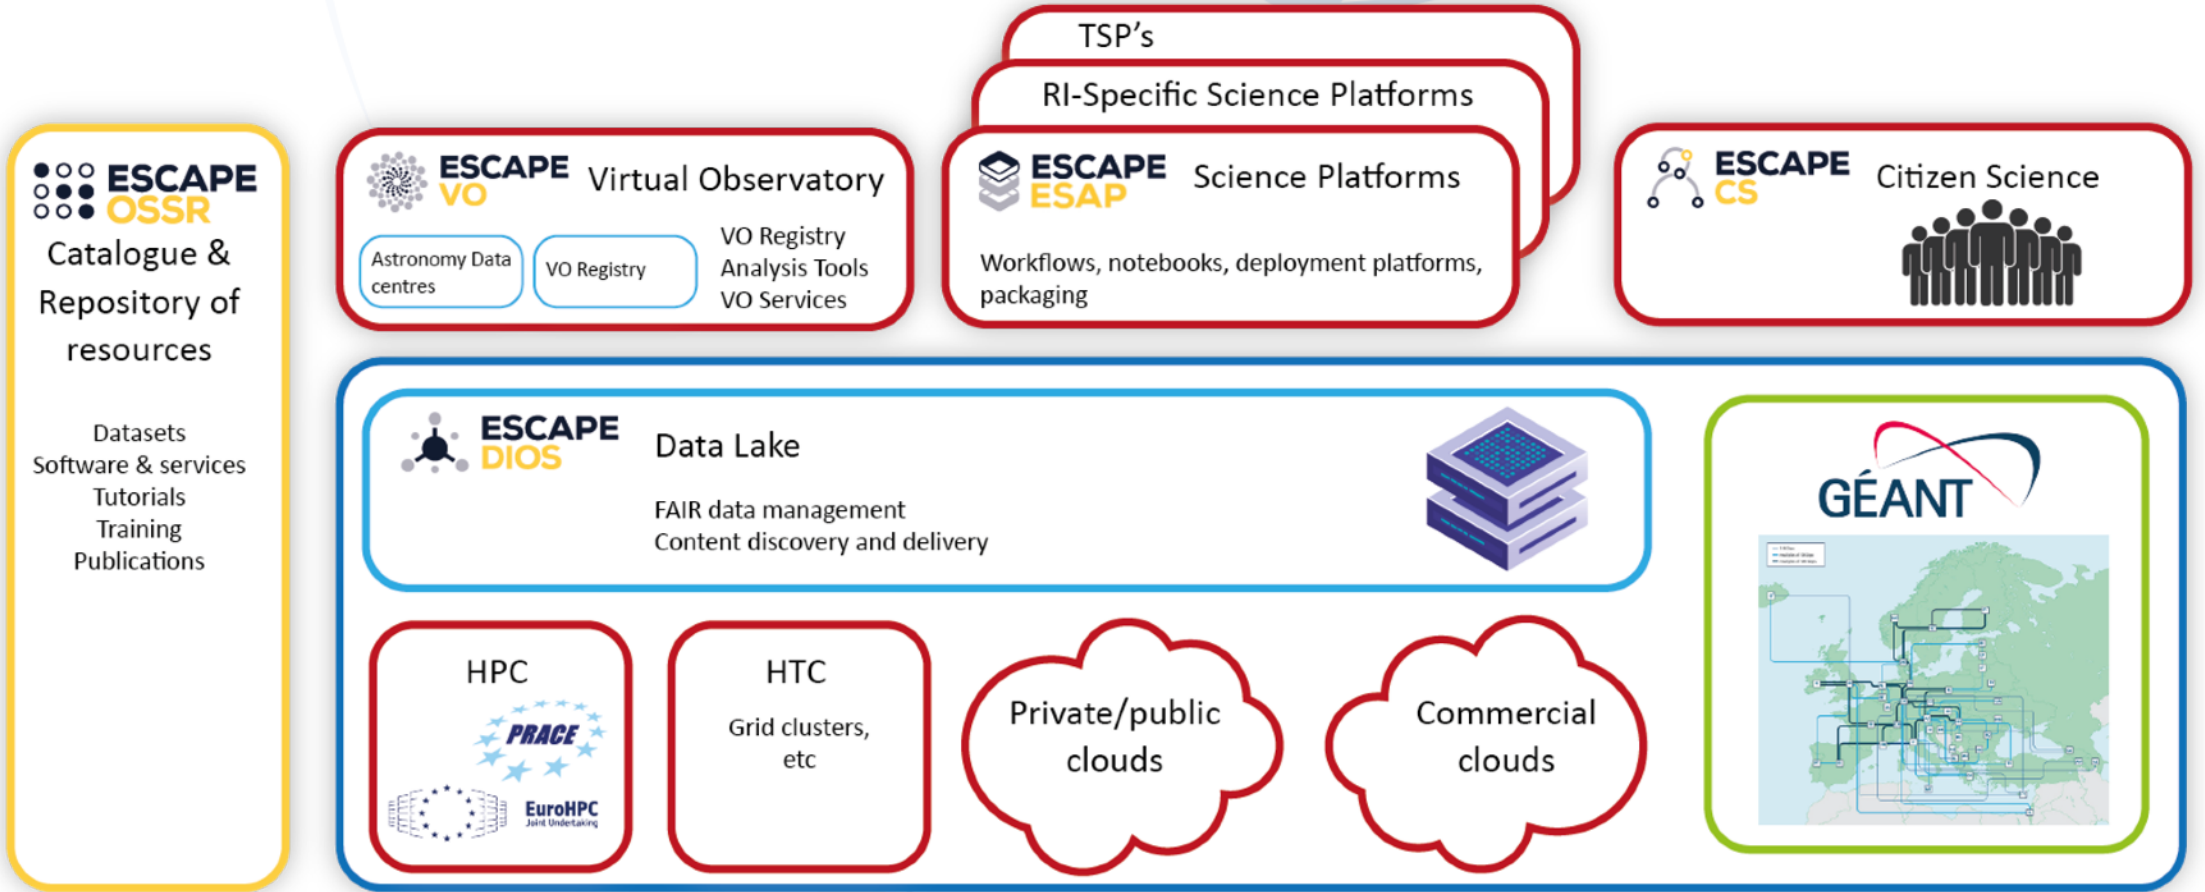
Acronyms

- **ESCAPE: European Science Cluster of Astronomy and Particle physics ESFRI** research infrastructures
- **ESFRI: European Strategy Forum on Research Infrastructures**
- **EOSC: European Open Science Cloud**
- **DIOS: Data Infrastructure for Open Science** (Yan's presentation)
- **ESAP: ESFRI Science Analysis Platform** (Klaas' presentation)

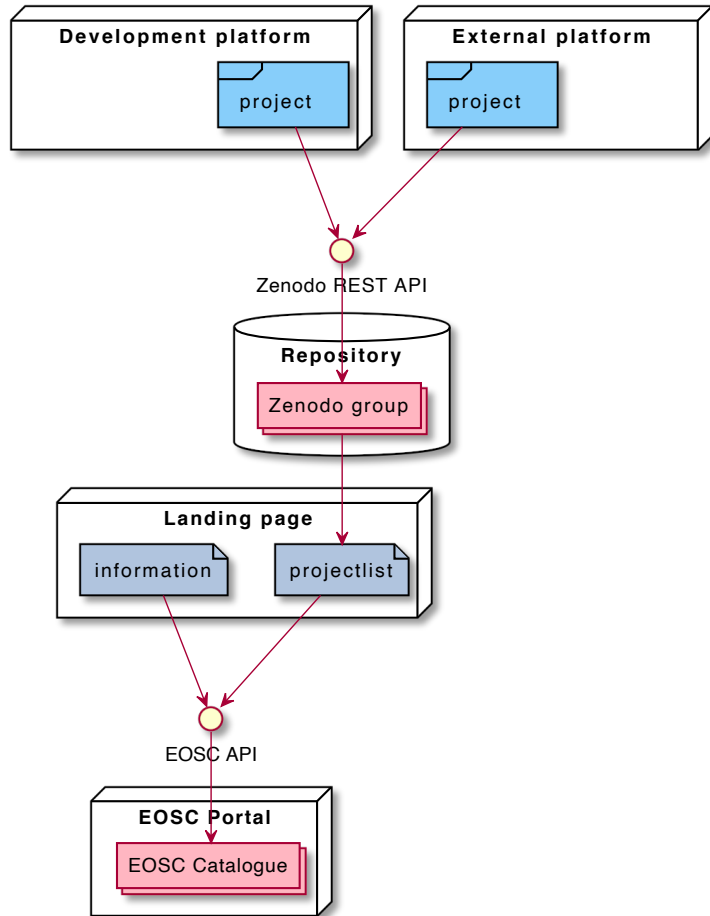
- So that means that “ESCAPE ESAP” is the European Science Cluster of Astronomy and Particle physics European Strategy Forum on Research Infrastructures research infrastructures European Strategy Forum on Research Infrastructures Science Analysis Platform.
- ...watch out for a quiz later!



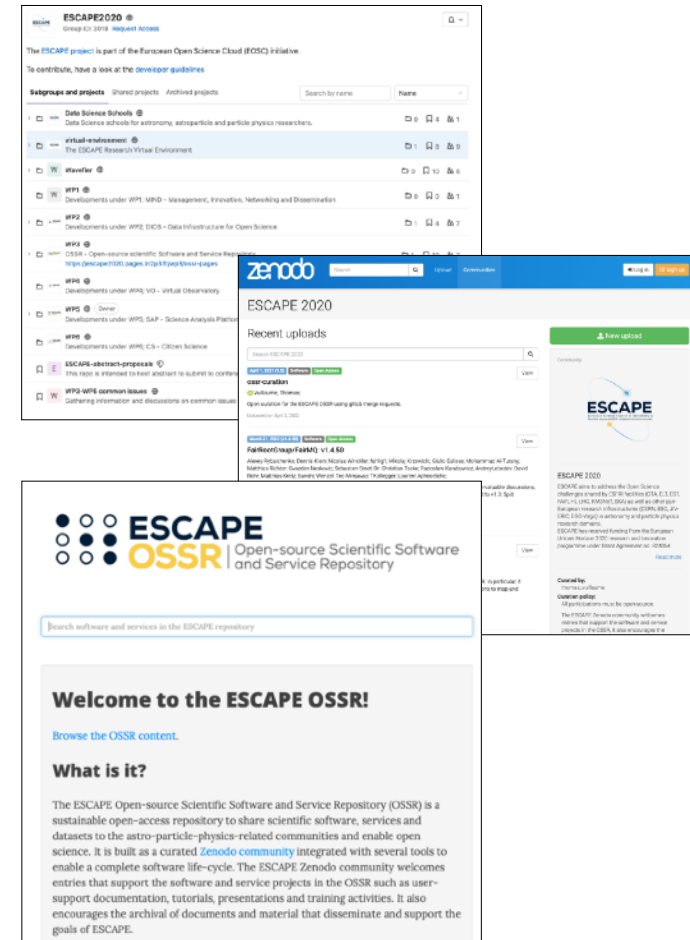
ESCAPE Work Packages



Software & Services Repository

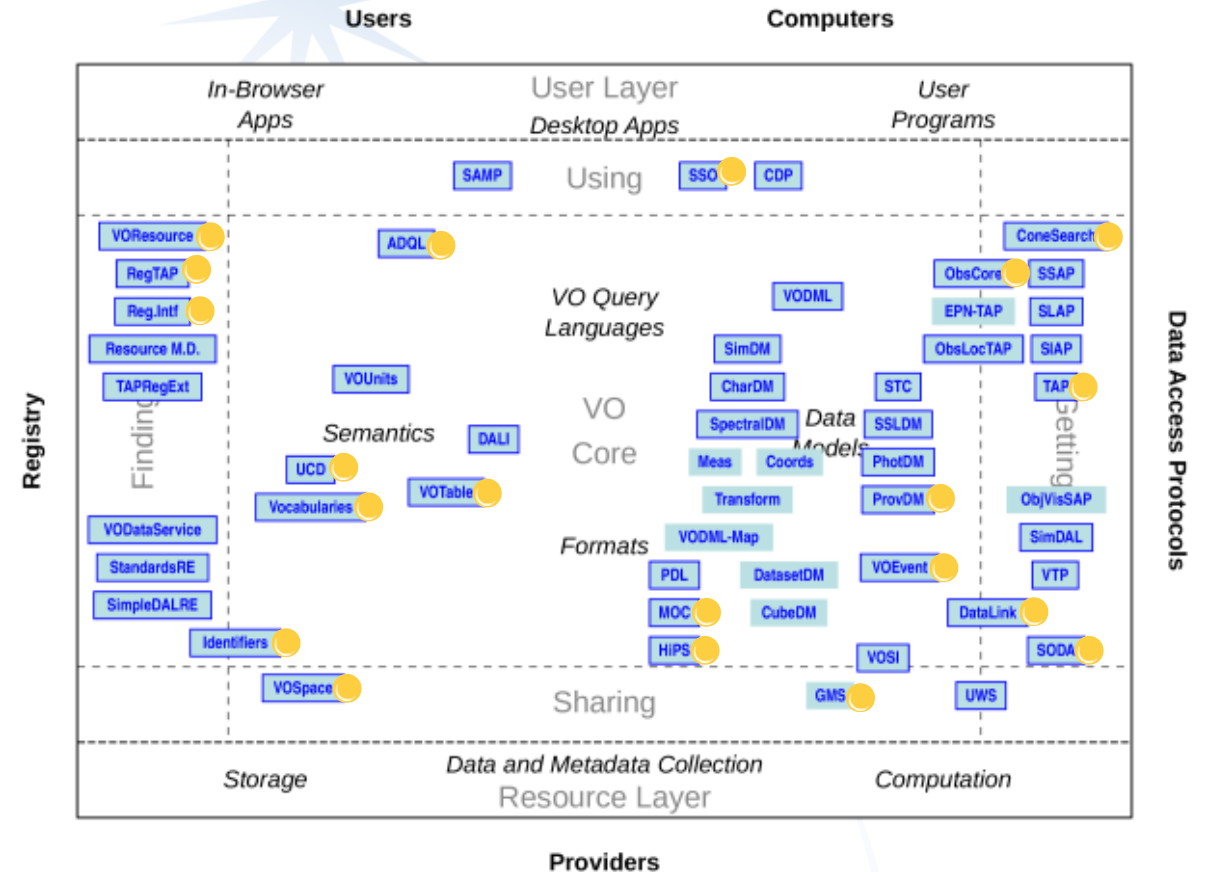


- A common repository for software across the various ESCAPE infrastructures.
- Designed both to *make software available* to practising scientists and to *preserve* software for future reproducibility.
- Built-in to common CI/CD systems (GitLab, Hub, etc).
- Built around the Zenodo system for long-term stability and preservation.
- Programmatic access through the **eOSSR** library.
- Integrated with the ESAP platform... see later.



Pushing on open standards: the Virtual Observatory

- ESCAPE aims to build around open standards wherever possible.
- We have used and contributed to a bunch of IVOA standards. In particular, the project has pushed on:
 - Using IVOA semantic UCD metadata for solar physics.
 - Developing IVOA standards for radio astronomy metadata.
 - Deploying IVOA-compliant services in the ESO science archive.
- Education, education, education, ...







The Scientific Data Lake

Yan Grange



Current situation

- Data sizes are increasing
- Distributed, federated storage
- Need to know what data is where
- Manual data management, database, ?
- Doesn't scale too well...



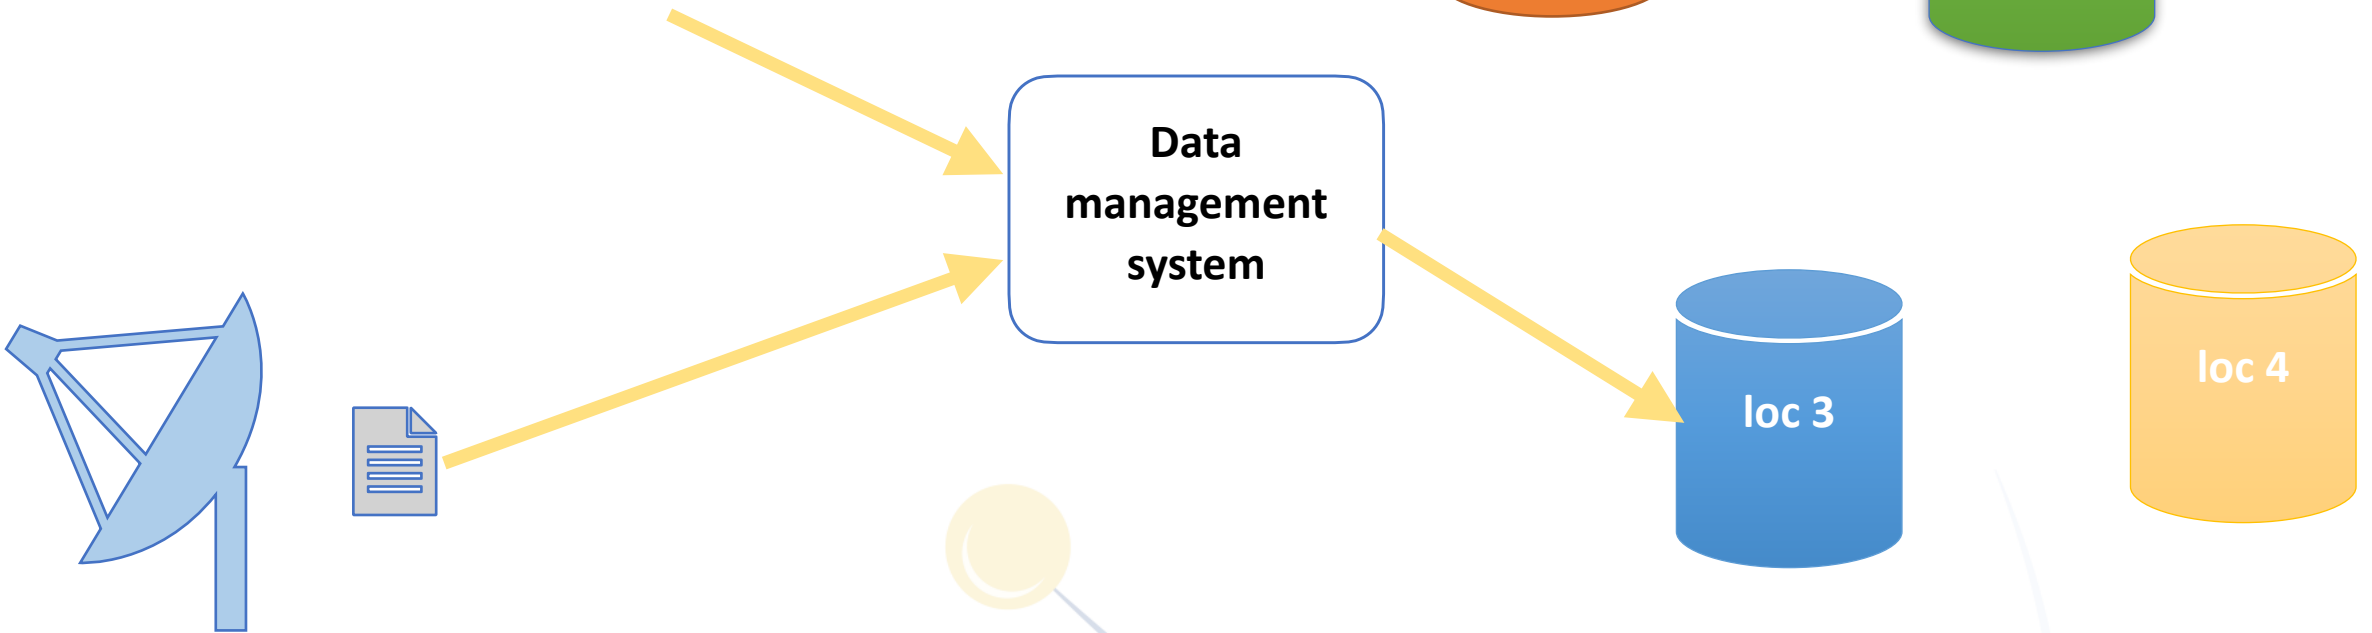
Scientific Data Lake (ESCAPE WP2)

- CERN experiments generate large data volumes, as does radio astronomy (especially SKA), and astronomy in general.
- In both cases, the data is provided to users using distributed infrastructure
- ATLAS experiment has developed an architecture for data management at this scale, which we will refer to as the Scientific Data Lake
- Of course, the domains do also have some significant differences; one of the goals of the DIOS WP is to investigate the architecture and find how it maps on the different domains, and feed it back to the Rucio development team.



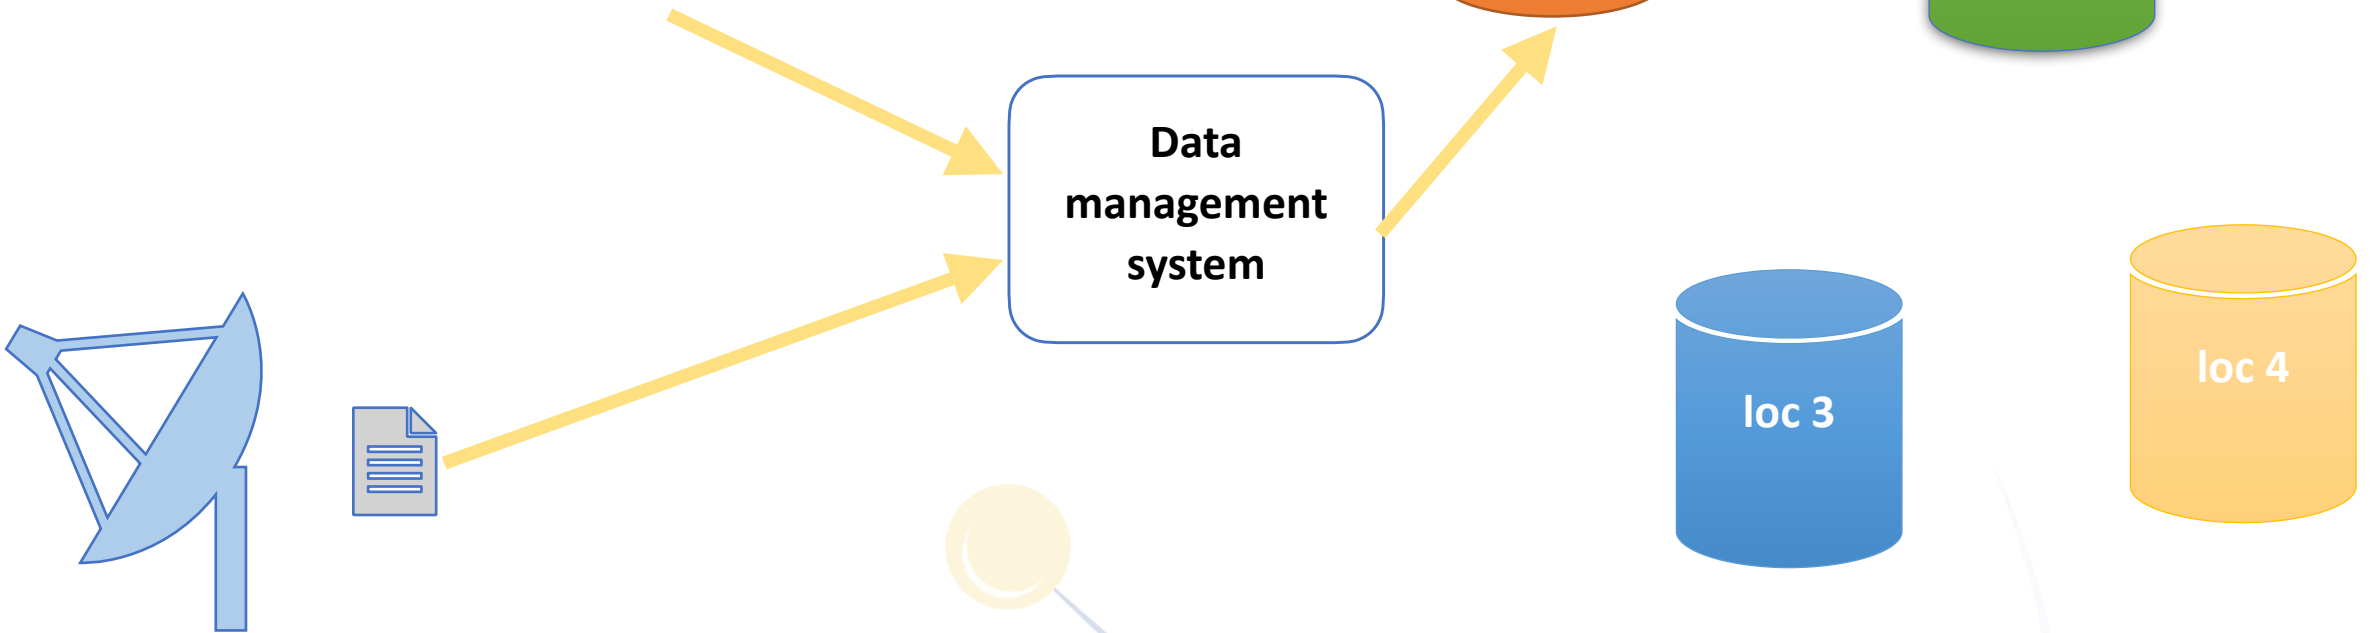
Rule-based data management

Survey observations go to blue



Rule-based data management


Pulsar observations go to orange



Rule-based data management

Two copies of FRB detection data

**Rules define the situation you want to achieve,
not (necessarily) how you achieve it.**





WIKIPEDIA
The Free Encyclopedia

[Main page](#)
[Contents](#)
[Current events](#)
[Random article](#)
[About Wikipedia](#)
[Contact us](#)
[Donate](#)

[Contribute](#)

Article [Talk](#) Read [Edit](#) [View history](#)  [More](#) Search Wikipedia 

Data lake

From Wikipedia, the free encyclopedia

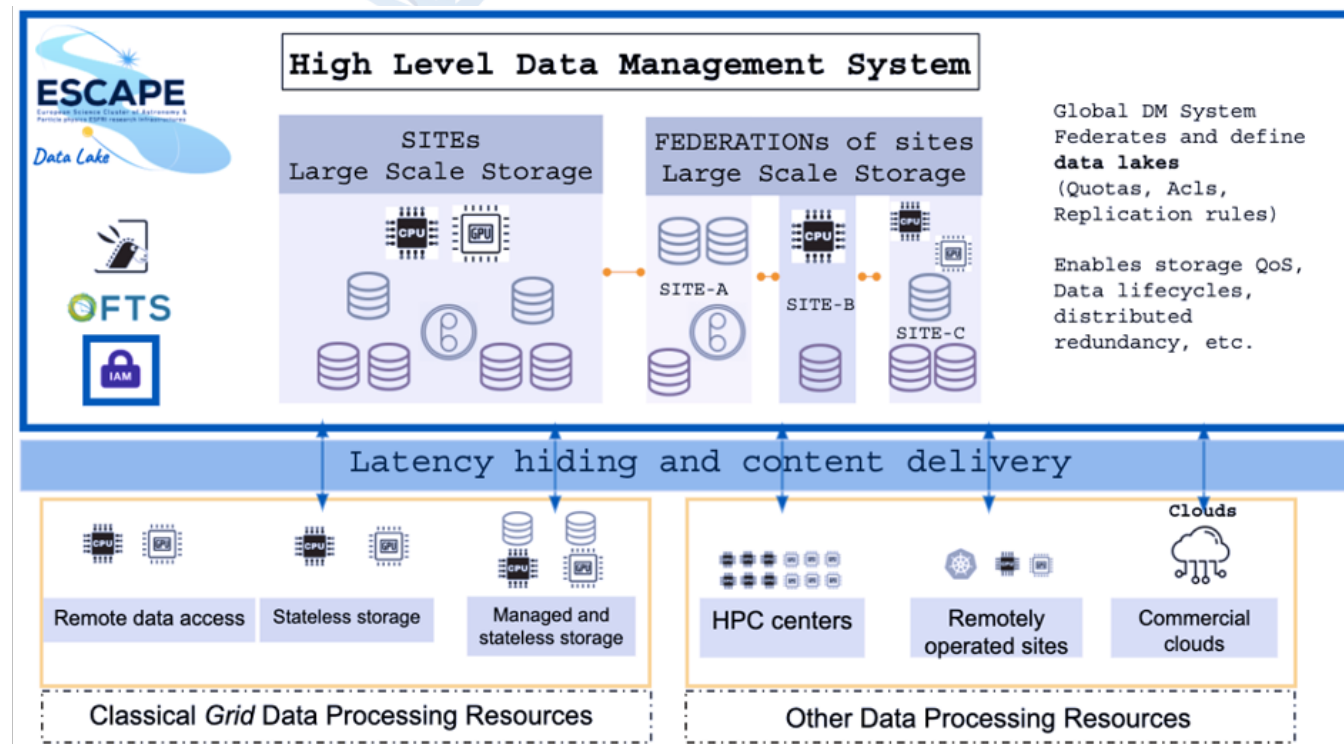
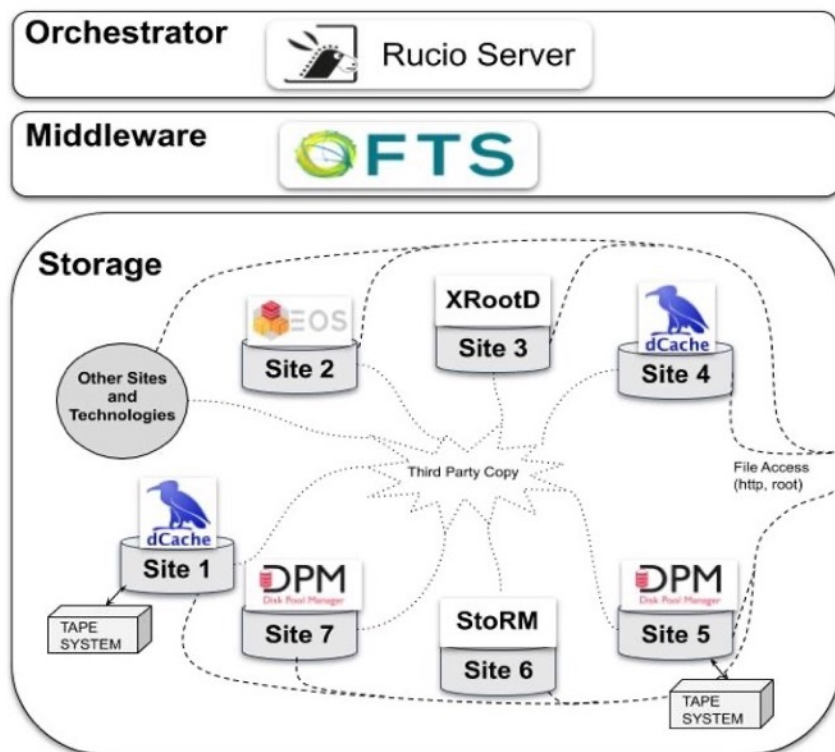
A **data lake** is a system or [repository of data](#) stored in its natural/raw format,^[1] usually object [blobs](#) or files. A data lake is usually a single store of data including raw copies of source system data, sensor data, social data etc.,^[2] and transformed data used for tasks such as [reporting](#), [visualization](#), [advanced analytics](#) and [machine learning](#). A data lake can include [structured data](#) from [relational databases](#) (rows and columns), semi-structured data ([CSV](#), logs, [XML](#), [JSON](#)), [unstructured data](#) ([emails](#), documents, [PDFs](#)) and [binary data](#) (images, [audio](#), video).^[3] A data lake can be established "on premises" (within an organization's data centers) or "in the cloud" (using cloud services from vendors such as [Amazon](#), [Microsoft](#), or [Google](#)).

Poorly managed data lakes have been facetiously called data swamps.^[4]

● Naming can be somewhat confusing. I may say Data Lake, but what I mean is this concept of a Scientific Data Lake.



Scientific Data Lake architecture



Data management and orchestration: **Rucio**
File transfer and data movement: **FTS**
Content delivery and latency hiding: **XCache**

Data Lake Information System: **CRIC**
AAI: **Indigo IAM**

INFN





From: Martin Barisits (2020)

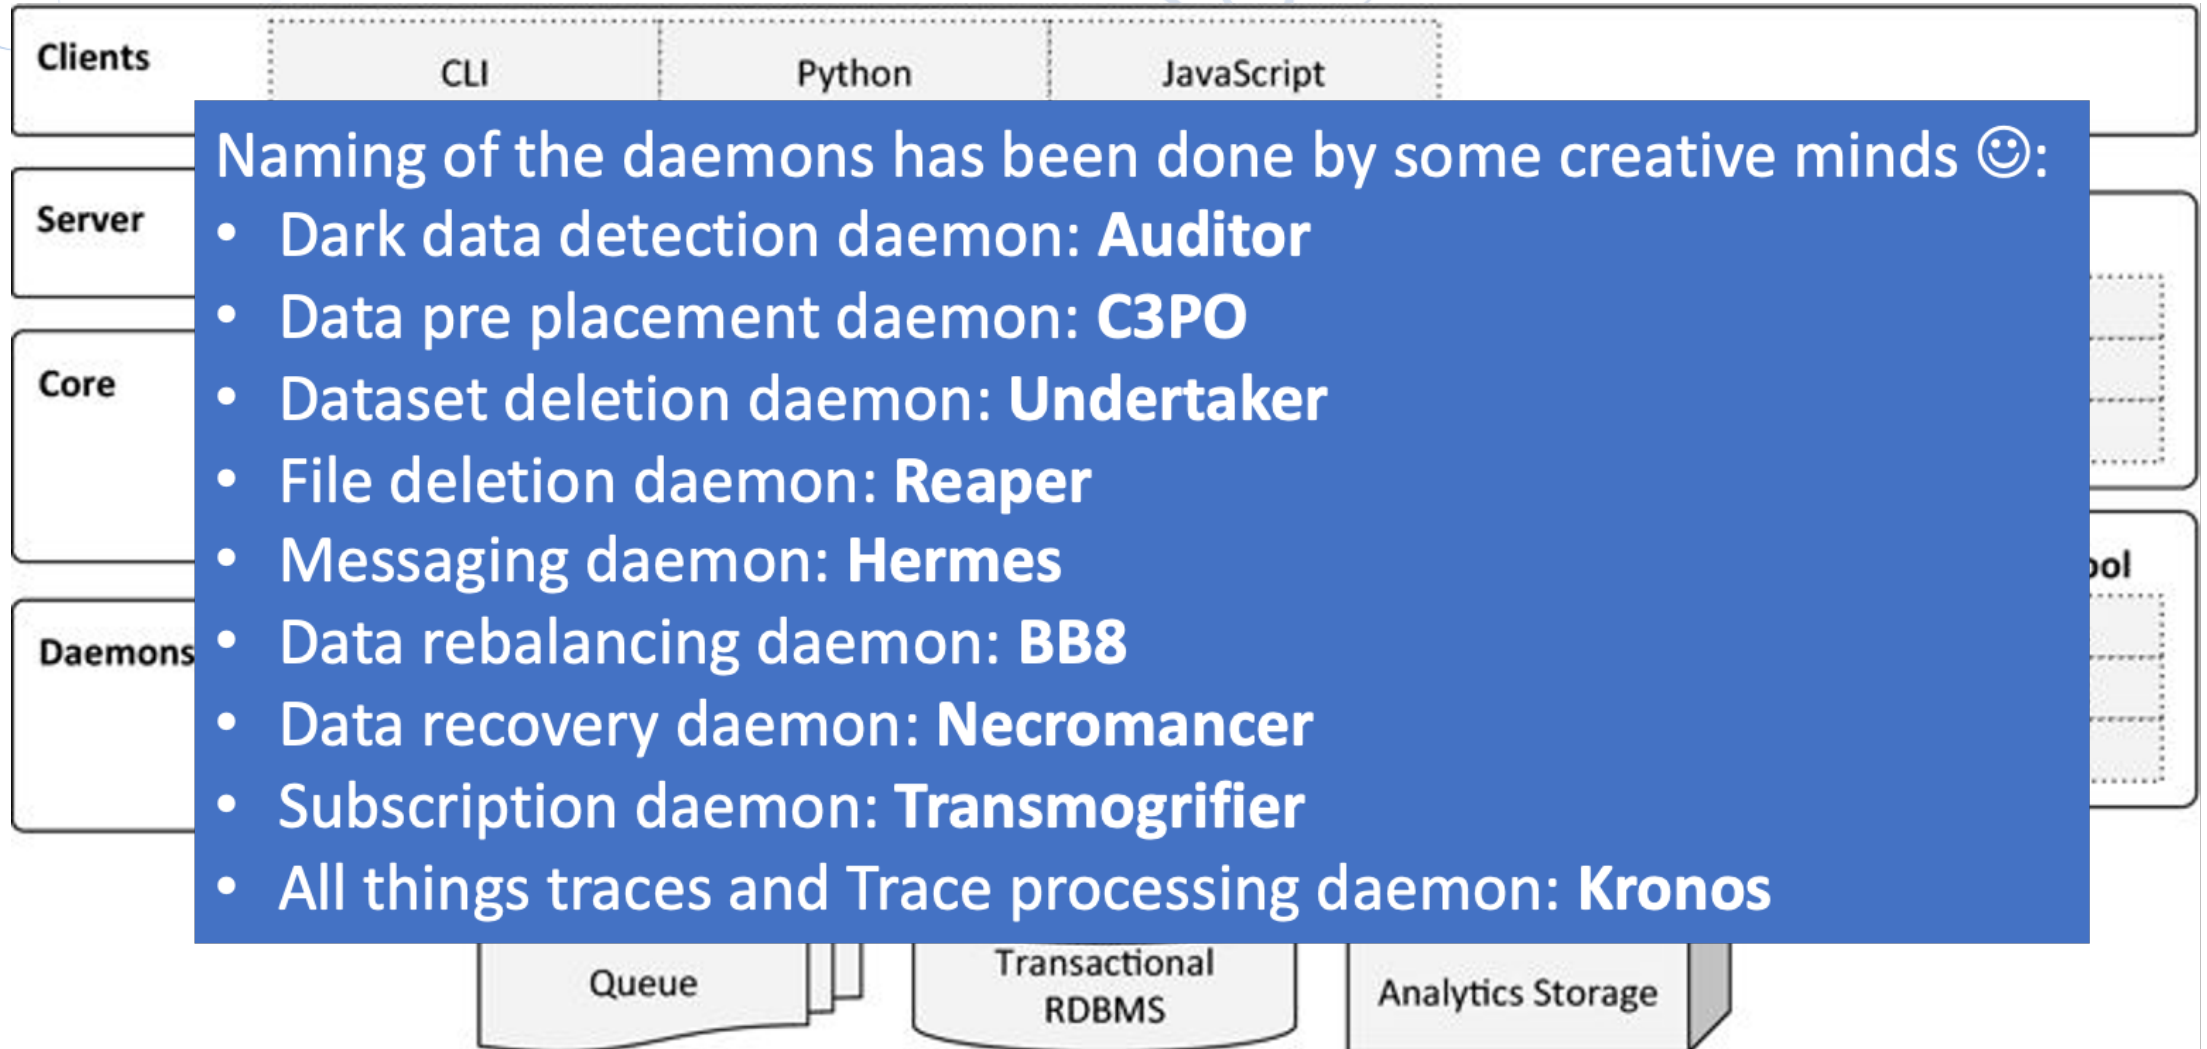


Rucio features

- Provides many features that can be enabled selectively
 - File and dataset catalog
 - Transfers between facilities including disk, tapes, clouds, HPCs
 - Web-UI, CLI, and API to discover/download/upload/transfer/annotate data
 - Extensive monitoring for all dataflows
 - Support for caches
 - Expressive policy engines with rules and subscriptions
 - Automated corruption identification and recovery
 - Data popularity based replication
 - ...

More advanced
↓





Rucio concepts

- Expressive metadata is key for rule-based management.
- Storage Elements are described by storage type (e.g tape, disk), space properties (e.g. used, free), geographical zone and supported protocol (e.g. https, s3, gridftp)
- Currently, some work is done in moving away from concrete storage device names (tape, disk, ssd) and move towards Quality of Service (fast, cheap, opportunistic) to support more expressive rules and not focus on the “how” but on the “what”.



Data in Rucio

- The metadata of the data is described using three types of attributes
 - System-defined attributes (e.g. file size, checksum)
 - Data attributes. By default those are particle-physics based (e.g. number of events, run number, task_id). But there is support for either JSON metadata or plugins.
 - Data management attributes (e.g. whether replicas should be purged if the data is not needed to be on a location anymore)



Rucio data hierarchy

Container

Alldata

Container

Data combination 1

Data set:

Observation 1

File 1a

File 2a

File 3a

Container

Data combination 2

Data set:

Observation 2

File 1b

File 2b

File 3b

Data set:

Observation 3

File 1c

File 2c

File 3c

- Name space of Rucio is flat. Everything (files, containers, data sets) is a **Data IDentifier (DID)**.

- Names are prepended by a “scope”

LOFAR_IMAGING:L12345_cal.MS.tar

scope name



Data management vs access

- Rucio supports data uploads and downloads, but primarily the goal of Rucio is data management (i.e. making sure data gets copied between remote end points)
- Manual uploads are certainly possible, but then you have to do the registration using either python or the REST API.
- Direct access to the data can be done either via a Download; downloading a DID which, as we saw before can be a file, or a set of files belonging together.
- Alternatively, data can be moved to a place that can be directly mounted (e.g. webdav, or fuse-mountable file system)
- Here one should keep in mind that by default, Rucio uses an algorithm to distribute files in directories. It can however be configured to behave otherwise.



Interacting with Rucio

- Primarily the Rucio daemon provides a REST API
- CLI (which in essence is a wrapper around the python library)
- Python library
 - This can be very neat if you want to integrate data management in your workflow, e.g. to manage the data ingests yourself.
- Web UI
- DataLake as a Service (which is a Jupyter notebook that gives access to a Scientific Data Lake; Klaas will show you more about that).



Interacting with Rucio

Rucio UI Monitoring Data Transfers (R2D2) Admin pattern OR name OR rule id Search Using account: grange

You are here: Data Identifier [LOFAR_ASTRON_GRANGE:L557202_SB004_uv.MS_09d196a3.tar] Rucio Version (WebUI / Server): 1.25.7 / 1.26.11

DID Metadata

account	grange
adler32	665847e7
availability	AVAILABLE
created_at	Tue, 17 Nov 2020 12:20:04 UTC
did_type	FILE
filesize	1.57 GB
guid	f6095af8ca49440ea0974f1a5d36db94
hidden	false
md5	22a83a04cce38601892cb8be3577e7d9
monotonic	false
name	L557202_SB004_uv.MS_09d196a3.tar
obsolete	false
purge_replicas	true
scope	LOFAR_ASTRON_GRANGE
suppressed	false
transient	false
updated_at	Sun, 13 Dec 2020 16:50:29 UTC

File Replica States

Show 10 entries Search:

Filename	Replicas
LOFAR_ASTRON_GRANGE:L557202_SB004_uv.MS_09d196a3.tar	EULAKE-1



Ongoing developments

- SKA has embedded a developer in the Rucio dev team, also SKA is prototyping with Rucio for the regional centre data management
- Transmogrifier daemon → making rules based on metadata
 - Implemented that this also works on custom metadata
 - Also implemented that this works with inequality
 - We could say: “put all data from the Galactic bulge in Amsterdam”
- Future: looking at performance of taking metadata from database (e.g. observation DB)



How to get started

- Requirement on FTS for data transfers. In principle several public FTS services are around, which could save people some head aches.
- One could point the Rucio server to one of those
- For the storage system behind it, use one that supports a protocol that FTS speaks (webdav is a safe choice to start with).
- Rucio is fully open source. There are helm charts for deployment on Kubernetes.
- For development and testing, also a development docker-compose file is present.



Some ESCAPE resources

- The ESCAPE team created docker containers (which can be turned in to Singularity) for the Rucio client, including the tooling that can be used to upload and download data (gfal). Of course the default configuration settings point to the ESCAPE instance, but this can easily be changed.
- <https://github.com/ESCAPE-WP2/Rucio-Client-Containers/tree/master/rucio-client-container>
- Also on the ESCAPE wiki there are some ‘data challenges’ that were executed by all the instruments
- e.g. the reports on <https://indico.in2p3.fr/event/22501/>



The ESFRI Science Analysis Platform

Klaas Kliffen



References

John Swinbank



References

- ESCAPE:
 - <https://www.projectescape.eu>
- Software Repository:
 - <https://escape2020.pages.in2p3.fr/wp3/ossr-pages/>
- ESAP Repositories:
 - <https://git.astron.nl/astron-sdc/esap-api-gateway>
 - <https://git.astron.nl/astron-sdc/esap-gui>

