

ONEDATA

OPEN DATA LICECYCLE MANAGEMENT WITH ONEDATA

Michał Orzechowski

ACC Cyfronet AGH



WHO WE ARE?



- 10+ years of devoted development - see github.com/onedata
- Open-source, developed at the **AGH University of Krakow** and **Cyfronet** data center
- We work tight with scientific communities on a case-by-case basic
- Our vision is to:
 - deliver a **data management** platform for large-scale and **distributed** problems,
 - address the challenges of global collaborative data sharing across **federated** organizational domains,
 - streamline data processing in **heterogeneous** data storage setups.
- Our funding comes from Polish and European grants and partnerships

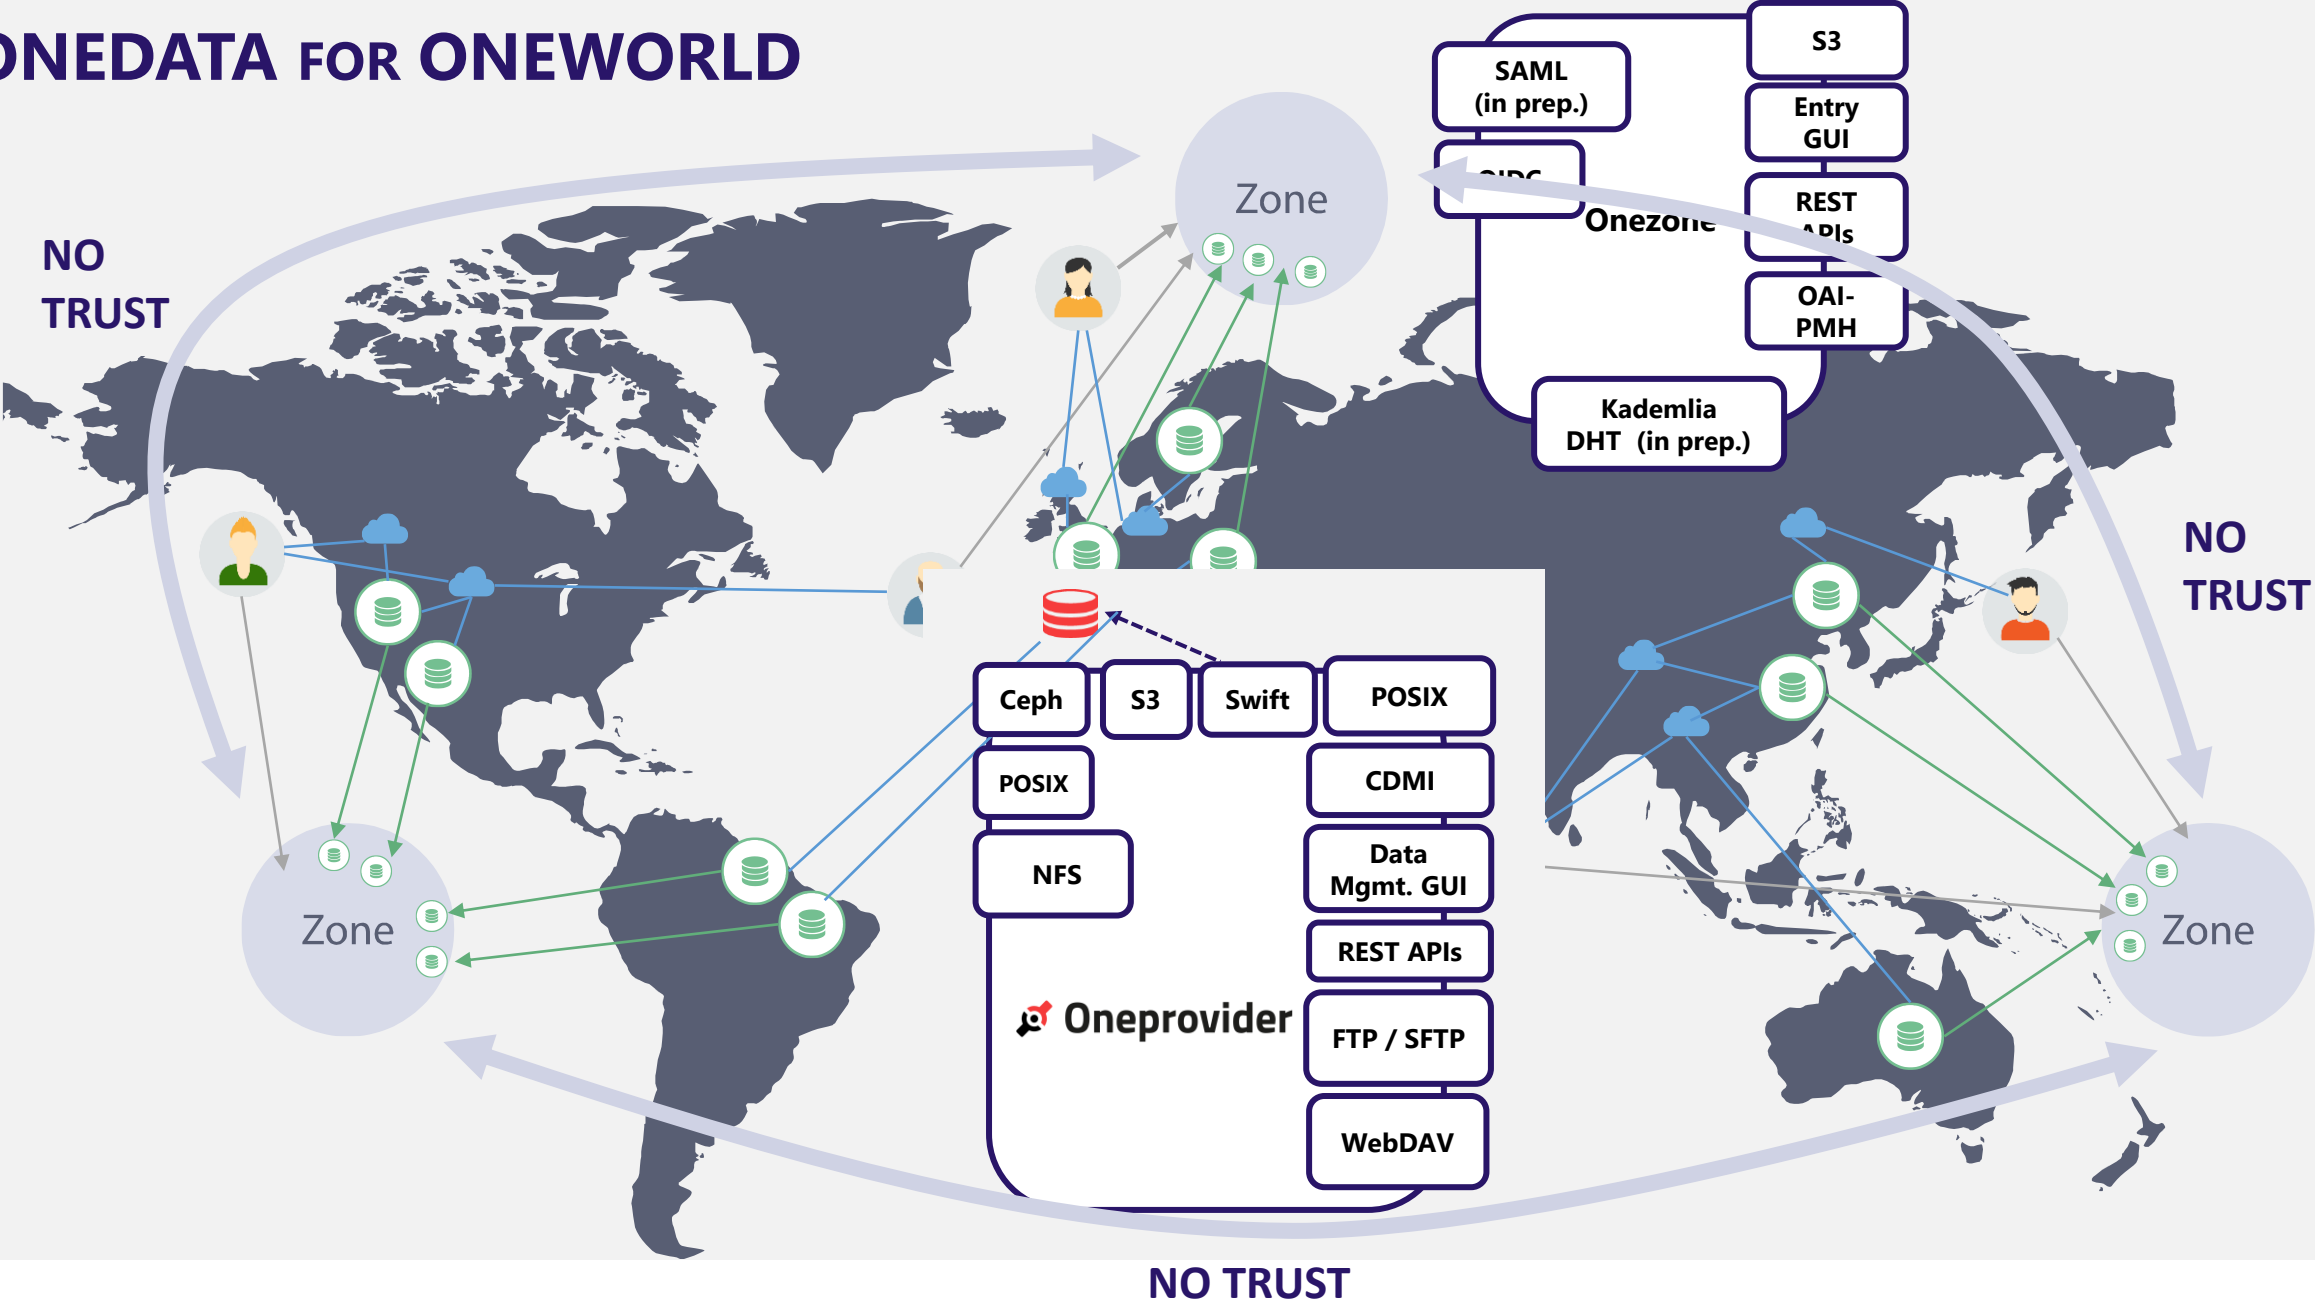
SUPPORTED BY SCIENTIFIC COMMUNITIES

- We are always looking for new partnerships and projects in order to:
 - keep the project running (of course),
 - gain invaluable experience cooperating with experts, solving real use cases, and working on authentic large datasets (big files & large number of small files)
- Our supporters and partnerships:



9 YEARS OF ONEEDATA @  CS³

ONEDATA FOR ONEWORLD




ONEDATA INSTANCES

- Between 5 and 10 active Zones in Poland and EU (depending on project lifecycles).
- Several instances not maintained by us.
- EGI DataHub (on the map), long haul project:
 - 20 sites (Oneproviders)
 - 2150 data spaces
 - ~1.77PB total storage size
 - 700+ users
- Archive for Polish National Museums:
 - 5PB of data — the current phase
 - 10PB of data — target scale
 - ~100M files



PROBLEMS ADDRESSED BY ONEDATA

- 1 Multi-protocol transparent access to data “[...] but we want POSIX”
- 2 Heterogeneity of storage technologies
- 3 Replica Management
- 4 Easy Data Sharing and publication (DOI)
- 5 Metadata Management Integrated with Data Management Platform
- 6 Flexible authentication and authorization
- 7 Easy integration using API with external services
- 8 High-throughput data processing
- 9 Access to Legacy Data Collections
- 10 Open Data Management



CS3 2017

@CS-nh8xn · 8 subscribers · 51 videos

More about this channel >

Subscribe

Home

Videos

Playlists

Community



onedata



Onedata - Eventually Consistent Virtual Filesystem for Multi-Cloud Infrastructures

CS3 2017 · 84 views · 7 years ago

Onedata - Eventually Consistent Virtual Filesystem for Multi-Cloud Infrastructures. Michal Orzechowski.

<https://www.youtube.com/watch?v=ia0osTTMVCs>

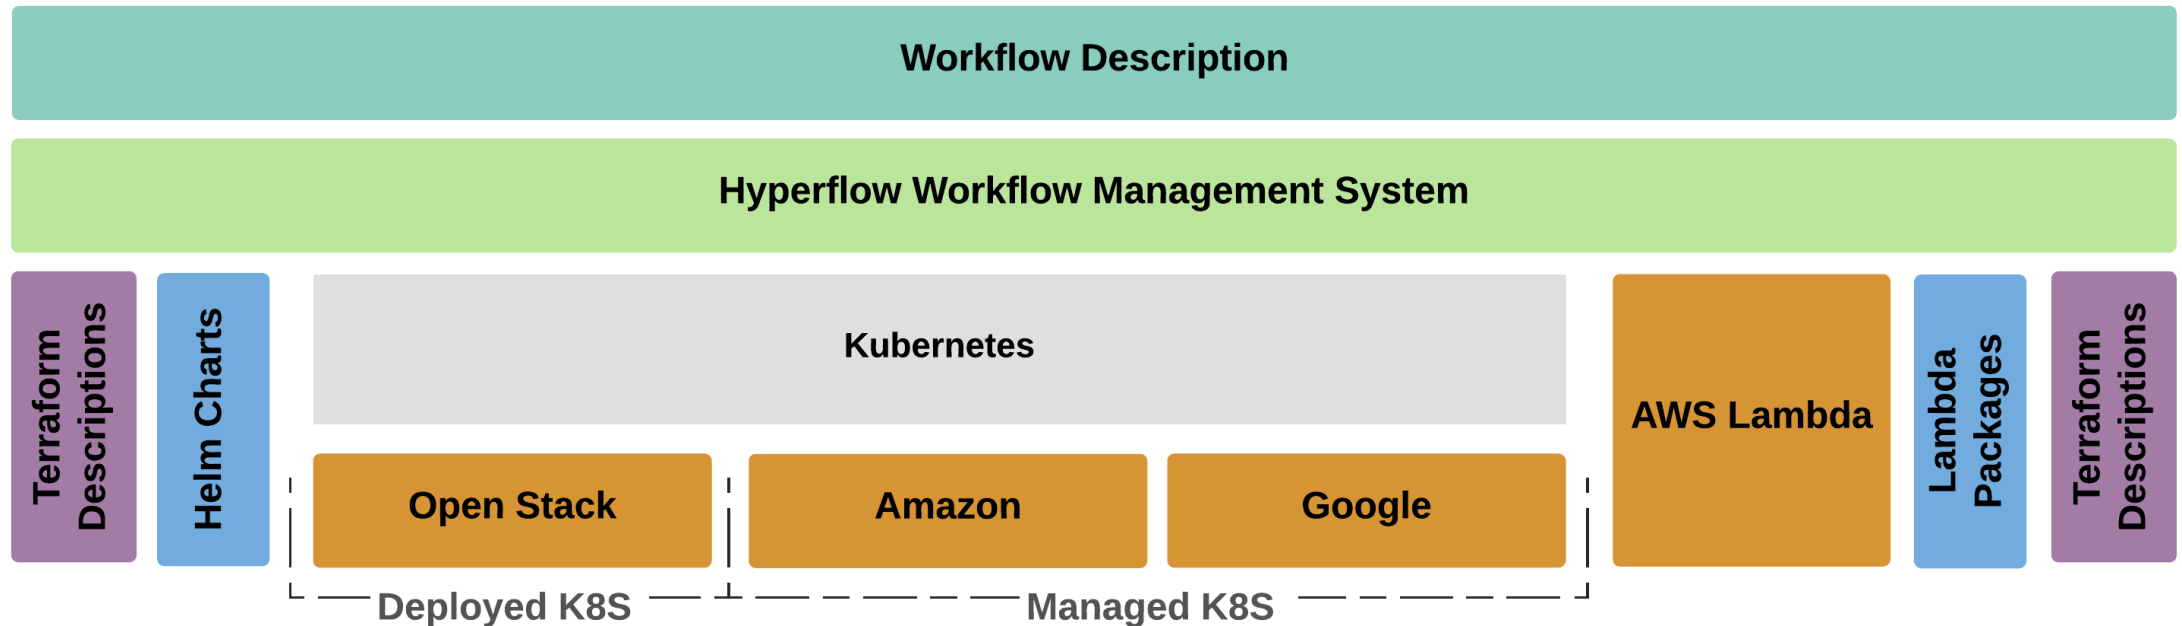
FIRST USECASE:

HYBRID-CLOUD SCIENTIFIC WORKFLOW EXECUTION

PROJECT REQUIRMENTS

Requirements:

- achieve automation using open-source DevOps tools (Ansible, Terraform etc.)
- use Kubernetes as a universal cloud API and resource manager
- use Hyperflow workflows manager schedules tasks to multiple Kubernetes clusters

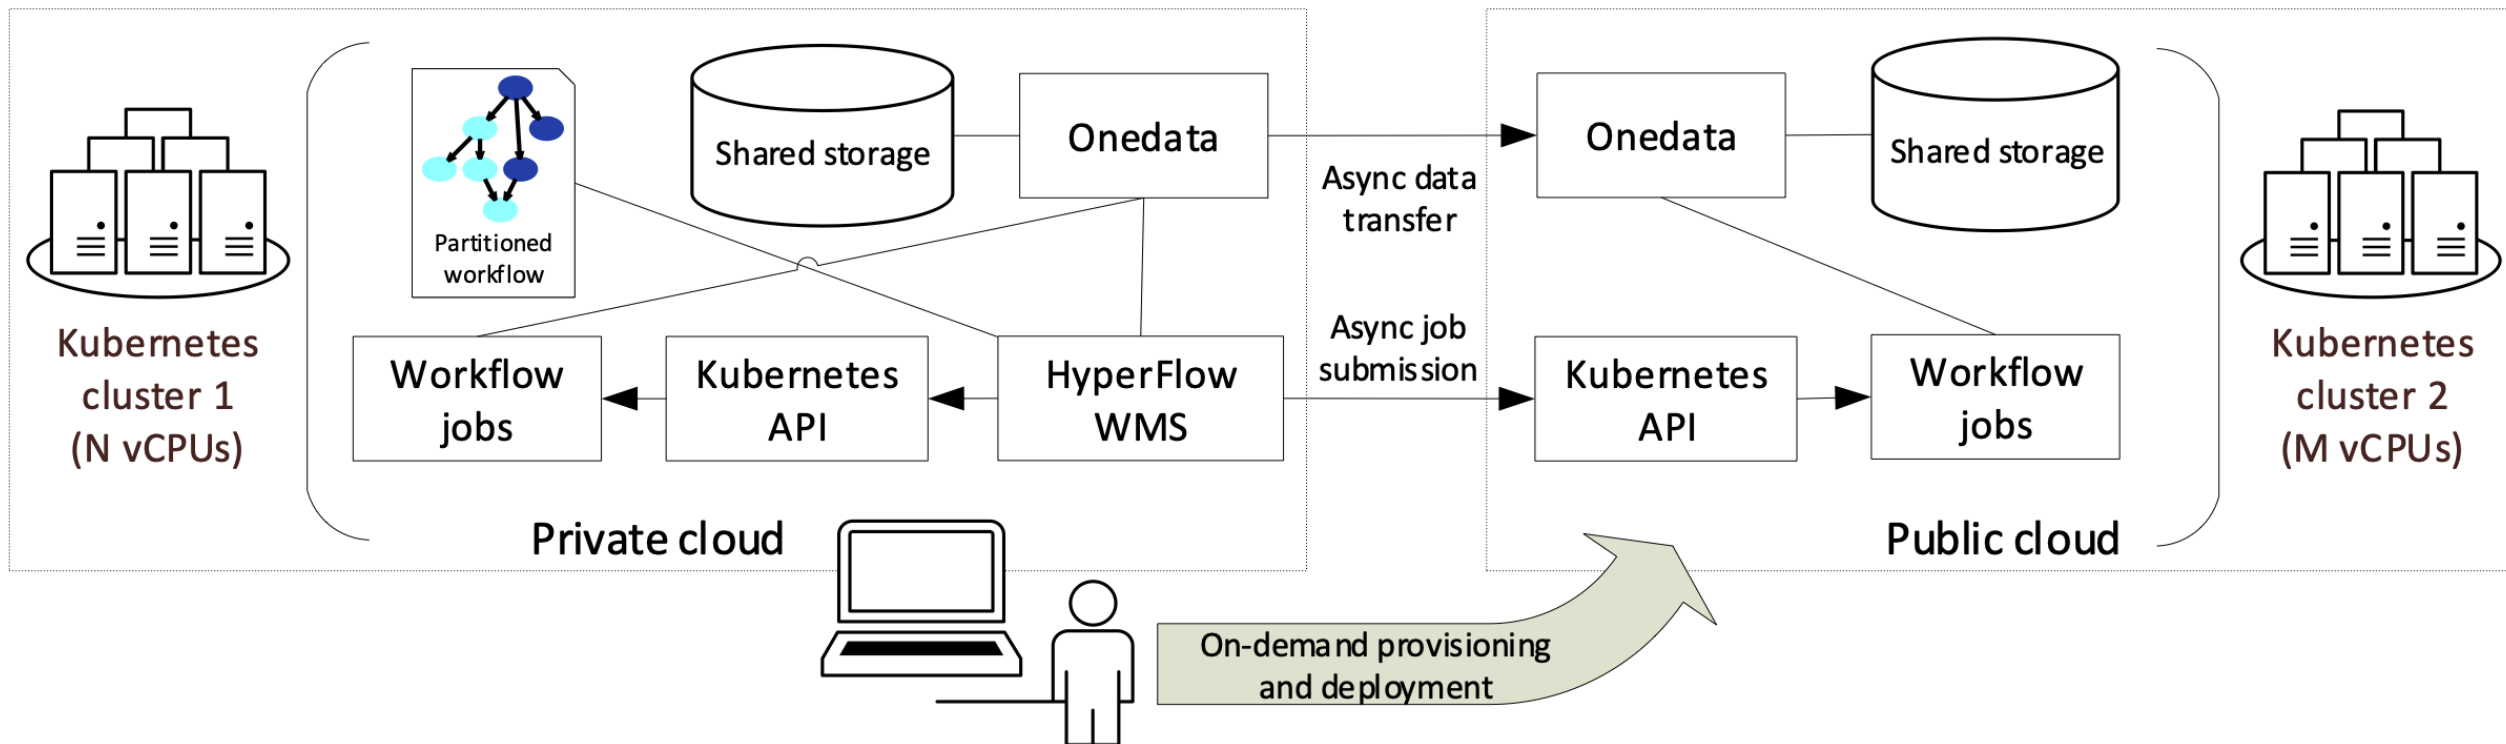


SCIENTIFIC WORKFLOW DATA MANAGEMENT

- designed for geographically-local filesystem (eg. Ceph, Lustre, NFS)
- required POSIX-compliant filesystem
- legacy software with unknown data access patterns (sequential/random)
- not designed for distributed data management or cloud-native storage systems

HYBRID-CLOUD DATA MANAGEMENT



- deployment of Hyperflow and Onedata and Kubernetes clusters
- workflow execution uses Onedata as a data-plane
- on-demand replication of data on the block level



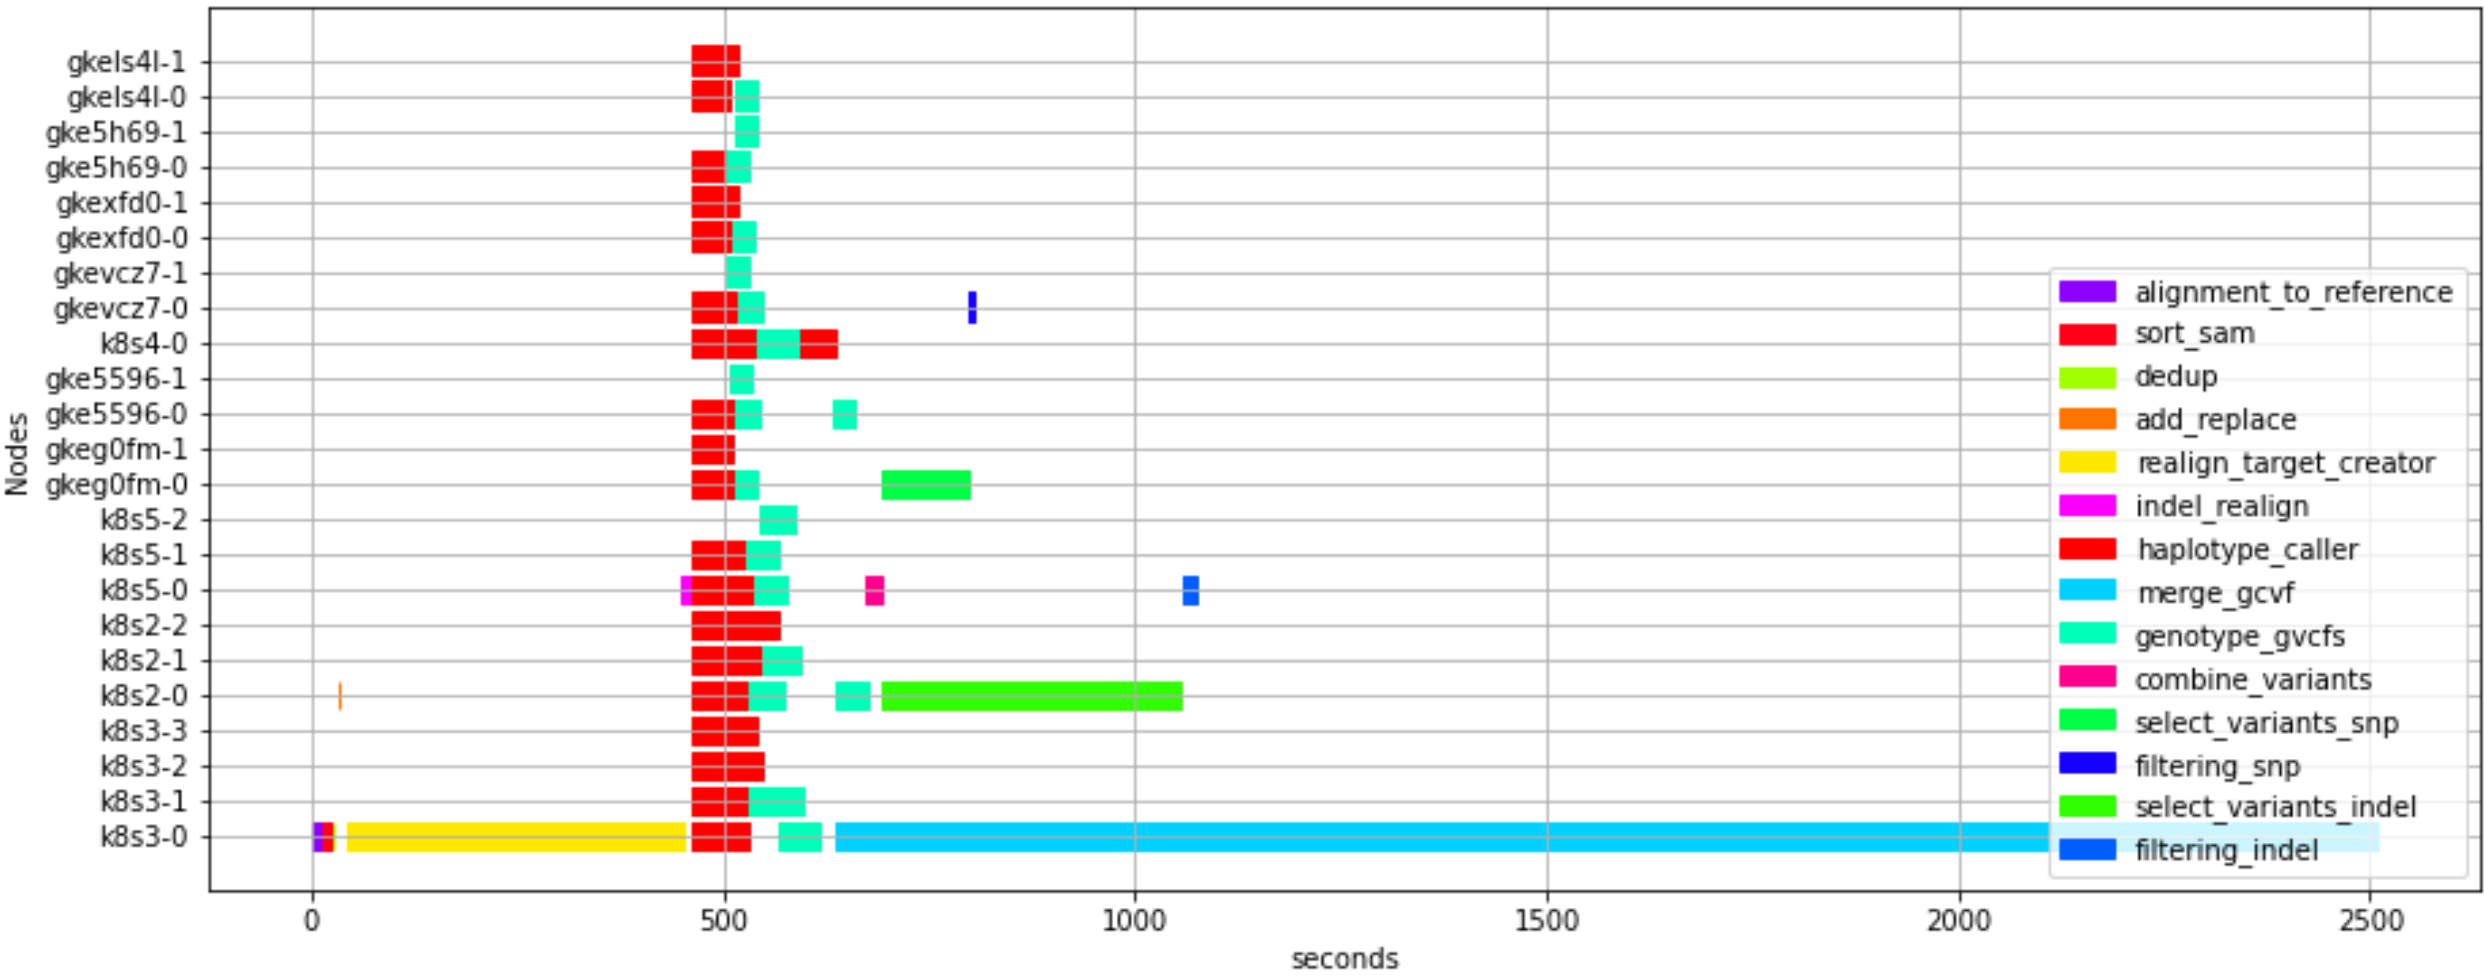
SOYKB GENOMIC WORKFLOW

Workflow for studying genomic variation of plants using next generation resequencing data with large-scale inputs, with following characteristics:

- a large input file – the reference genome database,
- two parallel stages (genotype_gvcfs and filtering_snp tasks),
- a long final task that merges the outputs of previous tasks (merge_gcvf)

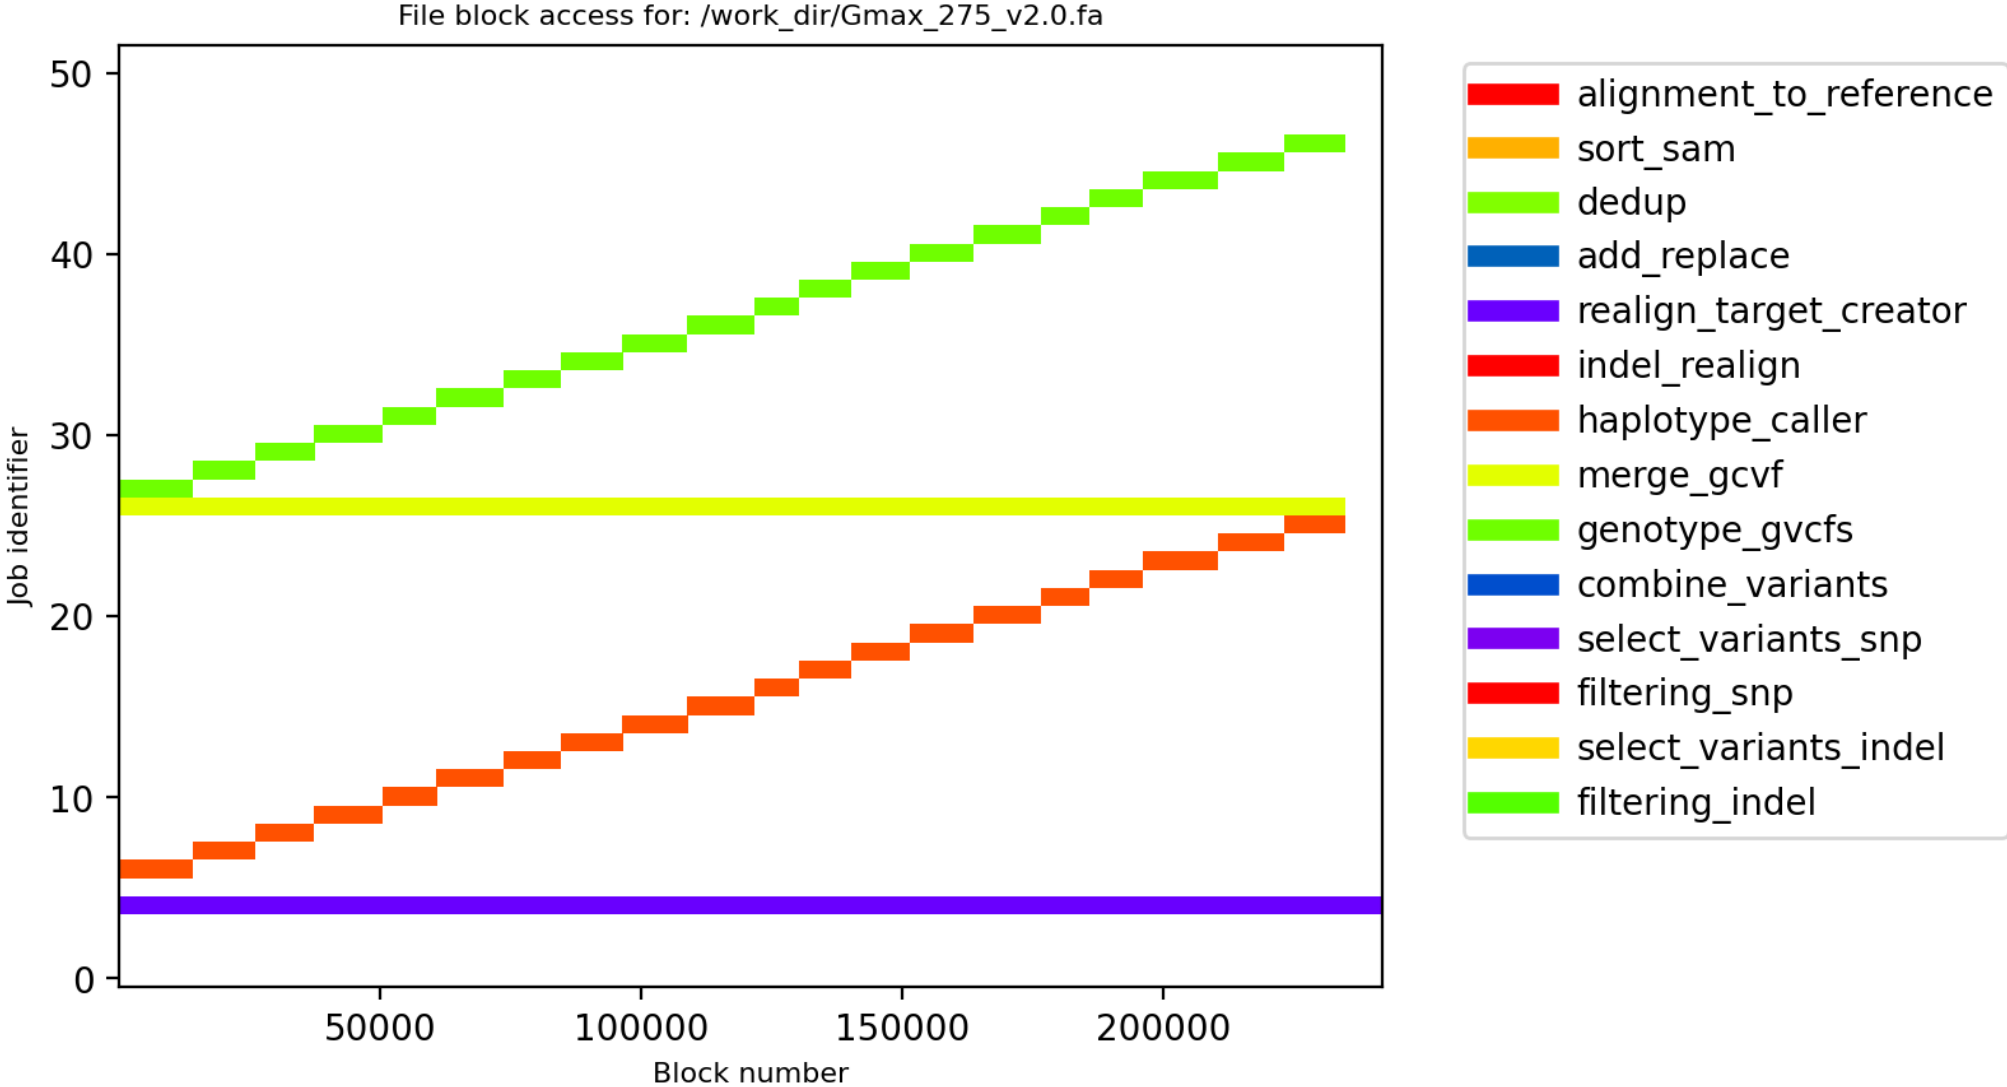


Structure of the Soykb workflow



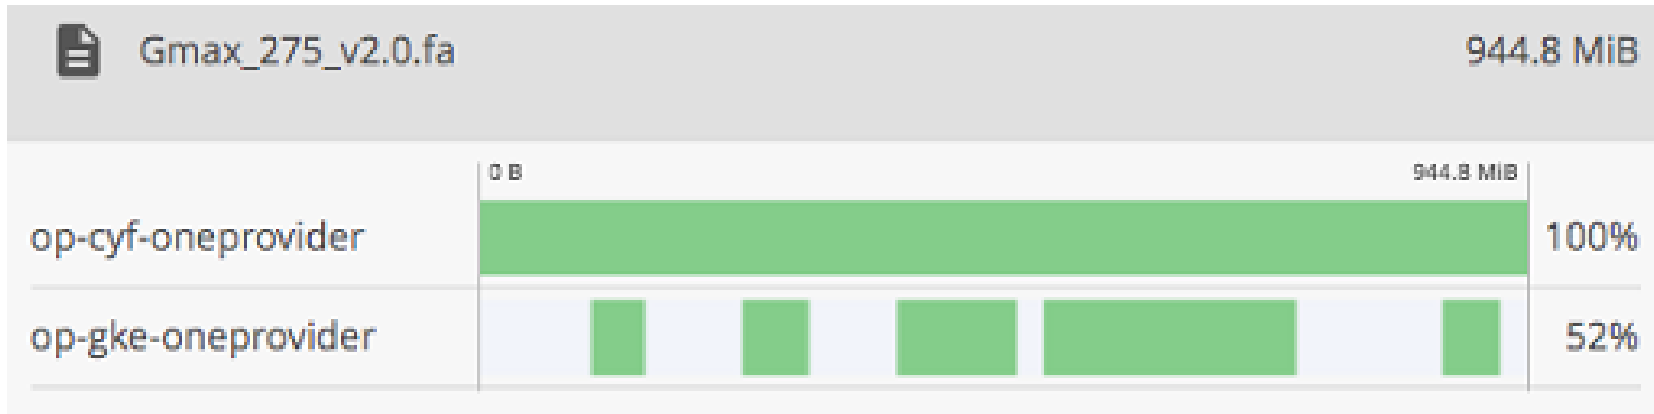
Visualization (trace) of the Soykb workflow execution on two clouds

SOYKB DATA ACCESS PATTERNS




Data access pattern of the largest input file of the Soykb workflow.

BLOCK LEVEL DATA DISTRIBUTION IN ONEDATA



Summarized data distribution of 165 files (2.5 GiB).

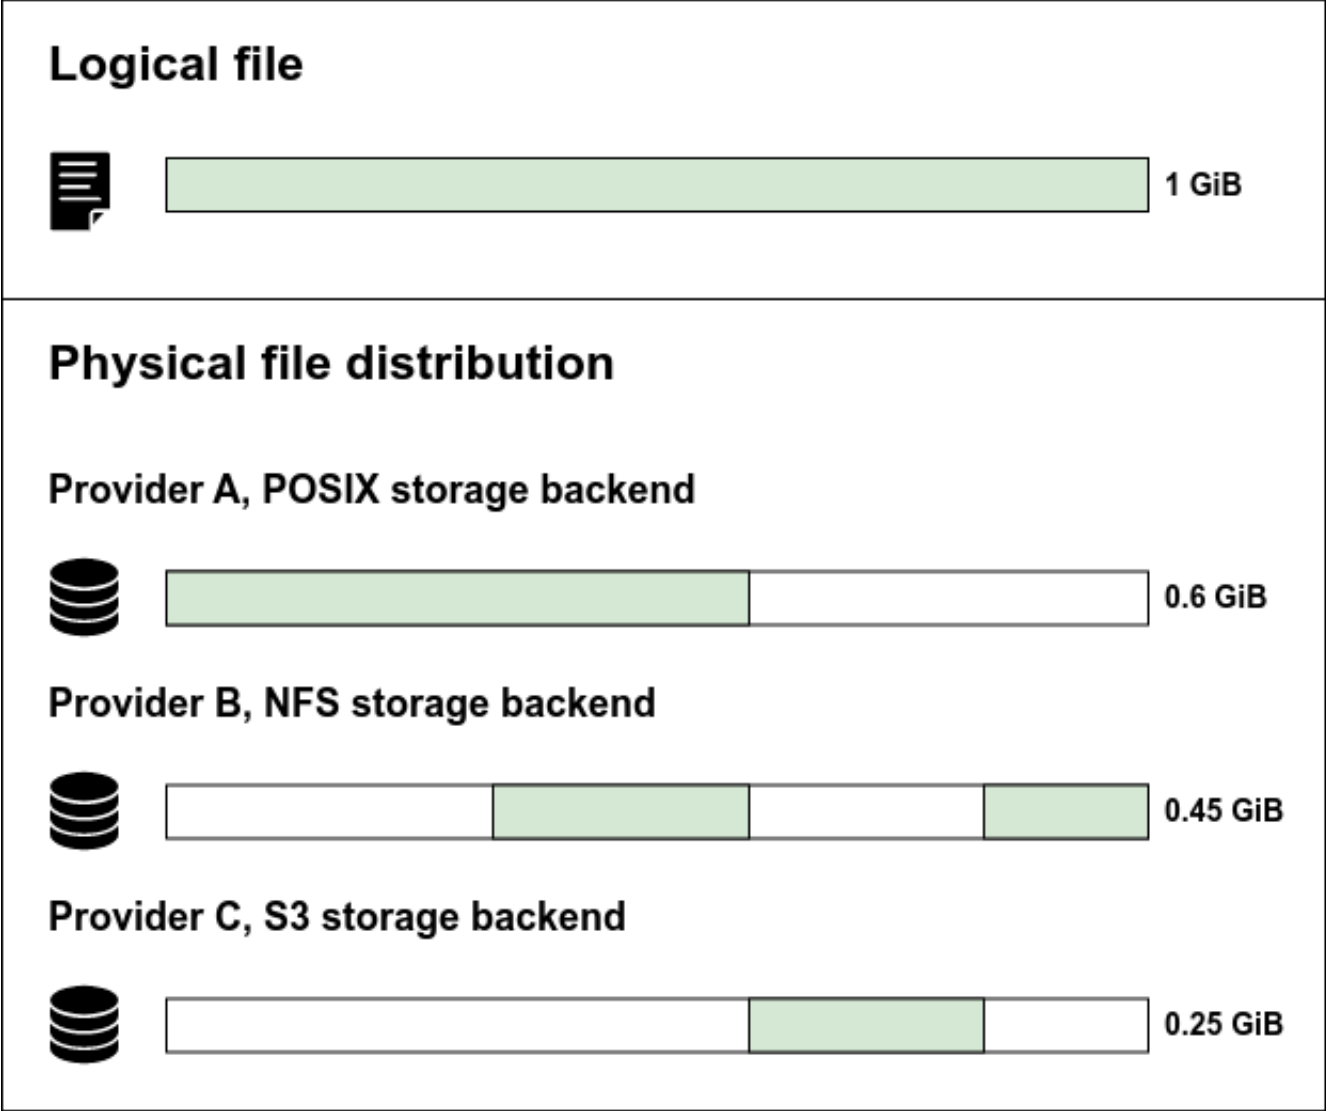


Runtime data distribution between two clouds for the Soykb workflow

DATA DISTRIBUTION

- the data in Onedata may be arbitrarily distributed among the storage backends of the supporting providers
- files are made up of parts of variable sizes — file blocks
- each provider holds a set of local file blocks, constituting a file replica
- when a file is read on a provider and the requested blocks are not present there, the missing ones are replicated on the fly from remote providers
- when a file is written on a provider, the overwritten blocks on other providers are invalidated. To read the file, the provider with invalidated blocks must once again replicate missing blocks from the provider with the newest version of the blocks

DATA DISTRIBUTION



FILE DETAILS



download-files.json

- Info
- API
- Metadata
- Permissions
- Shares
- QoS
- Distribution**

Data distribution per storage

onedatify @ **Bari**

/55eb75ab388522d56062f36fada12b40chb188/bw/download...



ceph @ **Krakow**

/55eb75ab388522d56062f36fada12b40chb188/a/2/7/a271c...




onedatify @ **Lisbon**

/55eb75ab388522d56062f36fada12b40chb188/bw/download...



posix @ **Paris**

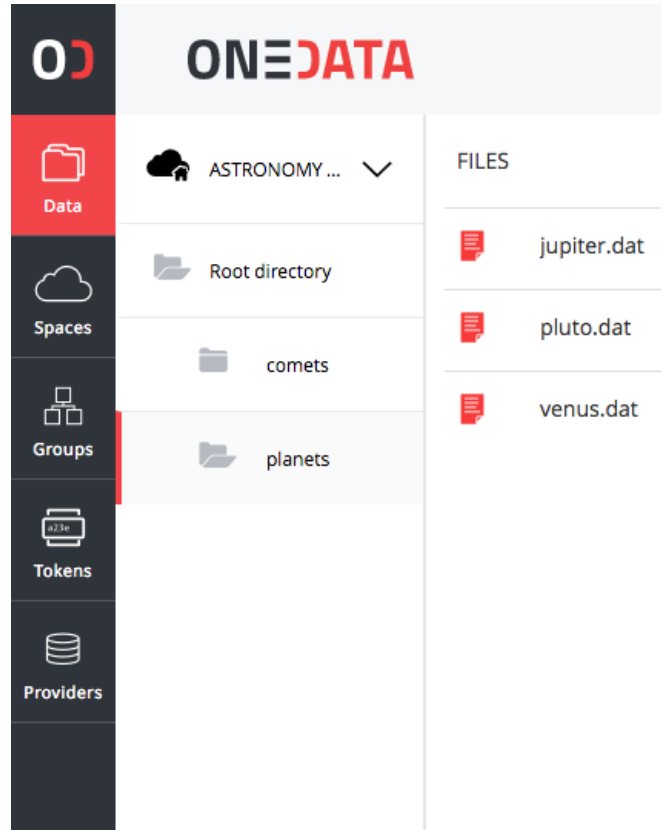


This file was transferred manually 1 time - [see history](#).

Block distribution

POSIX ACCESS WITH ONECLIENT

- presents Onedata virtual file system as POSIX
- support for most of the POSIX operations on globally distributed virtual file system
- all data accessible via a unified file system mountable on virtual machines, grid worker nodes and containers



```
[root@1f87c053280e oneclient]# ls
Astronomy Datasets  Big Data Experiment  Cancer Data
[root@1f87c053280e oneclient]# ls -lR
.:
total 0
drwxrwx--- 1 root 1733762 0 Sep 26 19:19 Astronomy Datasets
drwxrwx--- 1 root 1337123 0 Sep 26 19:14 Big Data Experiment
drwxrwx--- 1 root  608582 0 Sep 26 19:18 Cancer Data

./Astronomy Datasets:
total 0
drwxr-xr-x 1 1124656 1733762 0 Sep 26 19:20 comets
drwxr-xr-x 1 1124656 1733762 0 Sep 26 19:19 planets

./Astronomy Datasets/comets:
total 0
-rw-r--r-- 1 1124656 1733762 10000000 Sep 26 19:20 enck.dat
-rw-r--r-- 1 1124656 1733762 10000000 Sep 26 19:19 halley.dat

./Astronomy Datasets/planets:
total 0
-rw-r--r-- 1 1124656 1733762 10000000 Sep 26 19:07 jupiter.dat
-rw-r--r-- 1 1124656 1733762  5000000 Sep 26 19:08 pluto.dat
-rw-r--r-- 1 1124656 1733762  2000000 Sep 26 19:08 venus.dat

./Big Data Experiment:
total 0
-rw-r--r-- 1 1124656 1337123 10000000 Sep 26 19:08 cats_images.tgz
-rw-r--r-- 1 1124656 1337123  5000000 Sep 26 19:13 galaxies.img
-rw-r--r-- 1 1124656 1337123  5000000 Sep 26 19:14 spam_mails.tgz

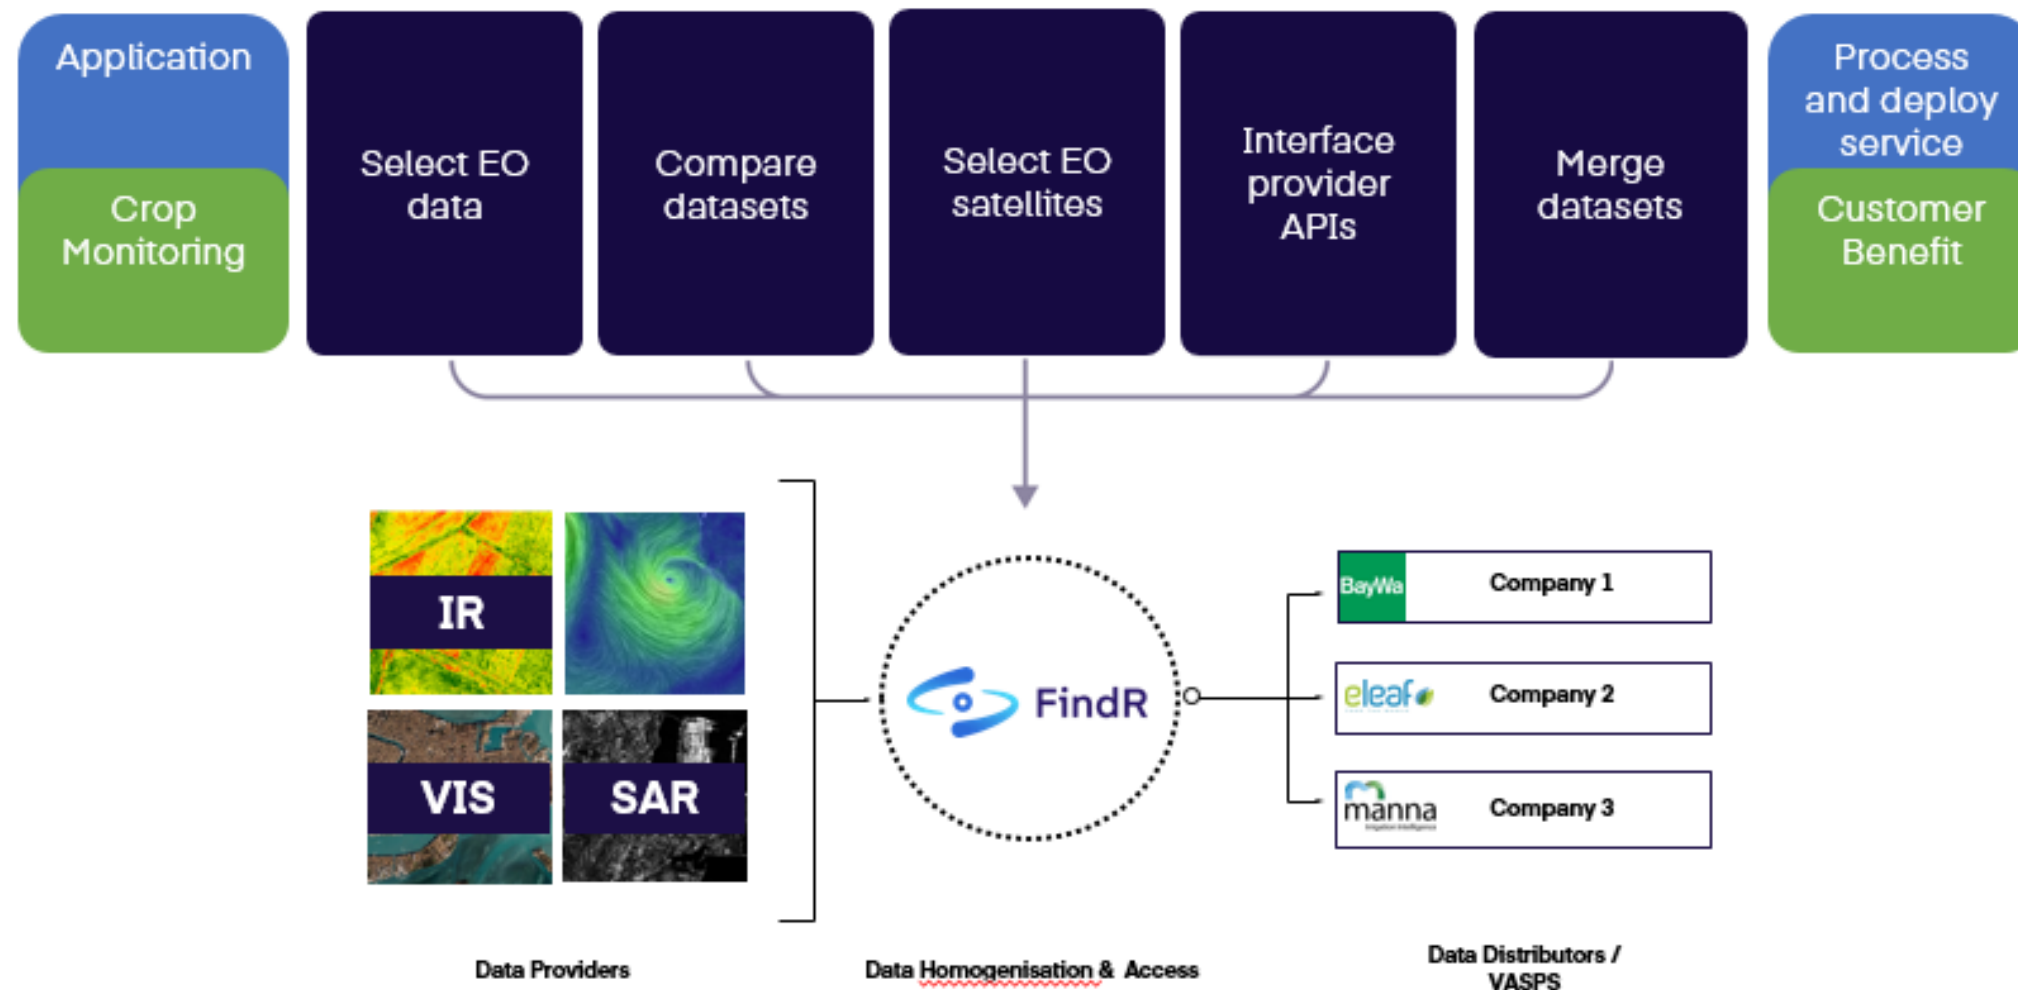
./Cancer Data:
total 0
-rw-r--r-- 1 1124656 608582 5000000 Sep 26 19:15 brain_tumor.zip
-rw-r--r-- 1 1124656 608582 5000000 Sep 26 19:14 duct_cancer.zip
[root@1f87c053280e oneclient]#
```

SECOND USECASE:

SATELITE IMAGE HOMOGENIZATION

SATELITE IMAGE HOMOGENIZATION

FindR provides **unified access** to several Data Providers and **homogenize** the data to a unified format. This includes a common **tiling scheme, resolution, spectral bands** and other characteristics of the output to enable the integration of new data sources without effort.

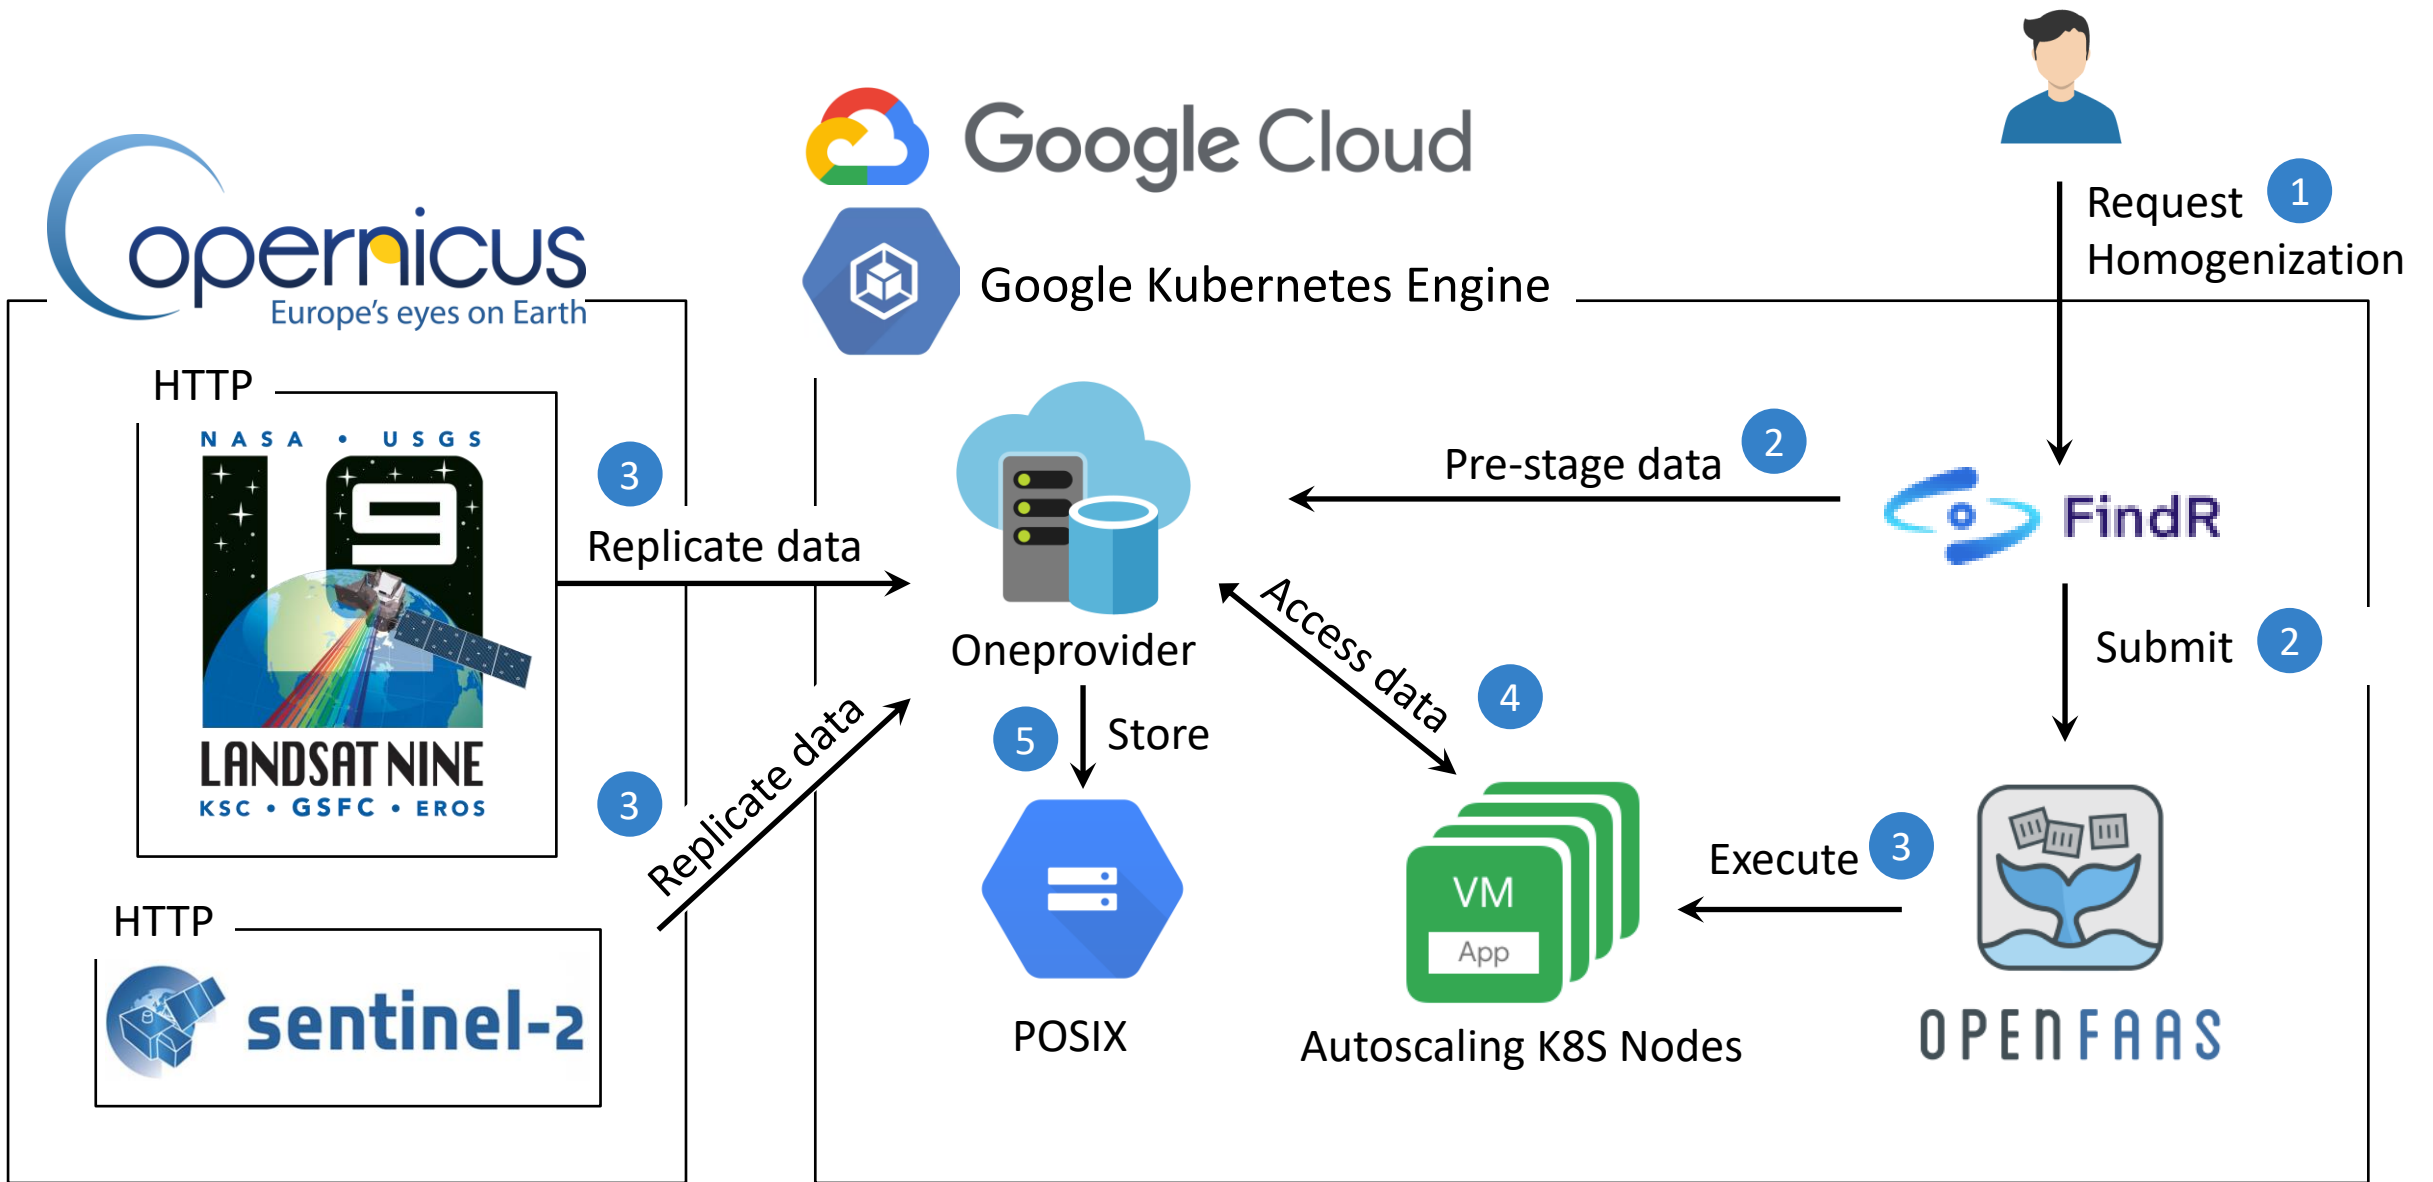


MULTIPLE STORAGE BACKENDS

Storage backends are used to store the physical data. Oneprovider accesses the storage backends via "helpers" (drivers) implemented for each supported type of storage. Helpers serve as a POSIX-like abstraction, building a layer over different storage backend APIs and access methods.

Currently supported storage backends:

- **POSIX** — any POSIX compatible filesystem accessible by Oneprovider via a mount point
- NFS — filesystem exported via the NFS protocol — no need to mount it locally
- S3 — Amazon S3 compatible storage
- Ceph RADOS — versions 14, 15, 16
- **HTTP** — any server exposing data via HTTP or HTTPS
- XRootD — CERN's data management protocol for LHC data
- WebDAV — experimental
- dCache — experimental



MANUAL DATA INDEXING

- Manual import allows registering specific files lying on a storage backend into a Onedata
- Preferred when only a part of the dataset is to be imported
- Upon registration, the file becomes accessible in Onedata under the requested path
- No data is copied in the process - a metadata entry is created in the space, containing a reference to the external storage.
- Manual storage import is supported on the following storage backends: POSIX, GlusterFS, WebDAV, XRootD, HTTP, S3, Swift, Ceph RADOS

```
curl -v -H "X-Auth-Token: $PROVIDER_ACCESS_TOKEN" -X POST "https://$PROVIDER_HOST/api/v3/onedata/provider/data/register" \
-H 'Content-Type: application/json' -d '{
  "spaceId": "'$MANUAL_IMPORT_SPACE_ID'",
  "storageId": "'$HTTP_STORAGE_ID'",
  "storageFileId": "https://packages.onedata.org/apt/ubuntu/2002/dists/focal/Release",
  "destinationPath": "/manually_imported_file"
}'
```

AUTO CLEANING & FILE POPULARITY

Overview Storage import File popularity **Auto-cleaning**


ENABLE AUTO-CLEANING

SELECTIVE CLEANING

- Lower size limit - Upper size limit - Not opened for - Opened up till now - Opened per hour - Opened per day - Opened per month

CLEANING BOUNDARIES

Total space: **1 TiB** Used space: **500 GiB** Free space: **524 GiB** To release: **248.1 GiB**



CLEANING REPORTS

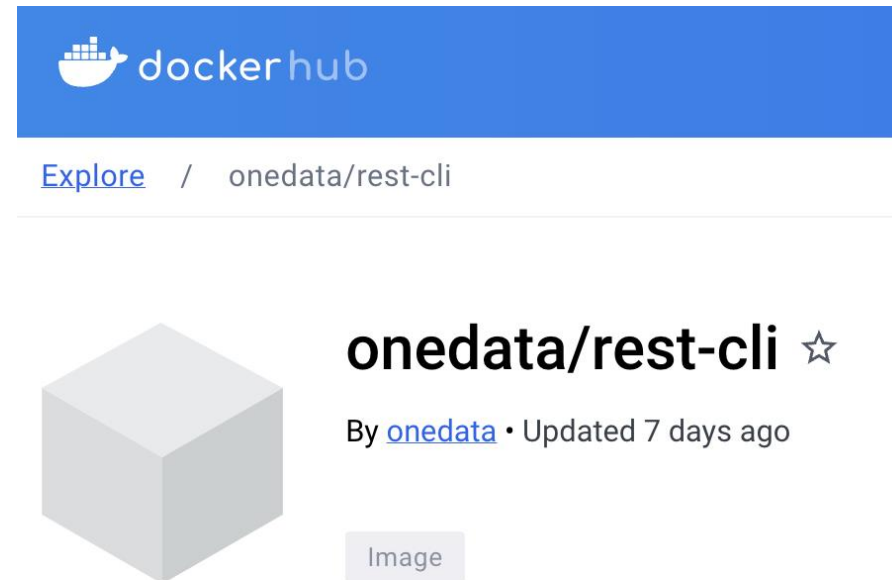
Start	Stop	Released size	Files number	Status
-------	------	---------------	--------------	--------

No reports available.

RICH COLLECTION OF APIS

Integrate external tools using rich API interfaces with data management platform and build more complex environments for data processing:

- APIs for all operations
- APIs for full eventually consistent integration with external systems
- API for data locality and replication
- Flexible permission checking for APIs
- API fully described using Swagger for generation of clients based on API specification
- Easy to use simple command line clients for REST API



**THIRD USECASE:
PROCESING OF LARGE MUSEAL DATASETS**



NATIONAL DIGITAL ARCHIVE



**BIBLIOTEKA
NARODOWA**

NATIONAL LIBRARY



Muzeum Narodowe
w Krakowie

NATIONAL MUSEUM IN
KRAKOW

	NATIONAL DIGITAL ARCHIVE	NATIONAL LIBRARY	NATIONAL MUSEUM IN KRAKOW
Number of files	4 millions	3 millions	37 millions
Number of dirs	1	33667	154578
Size	2068 TB	960 TB	13 TB

In Total: 44 million files and 3 Petabytes of data.

Planned: 100 million files and 10 Petabytes of data.

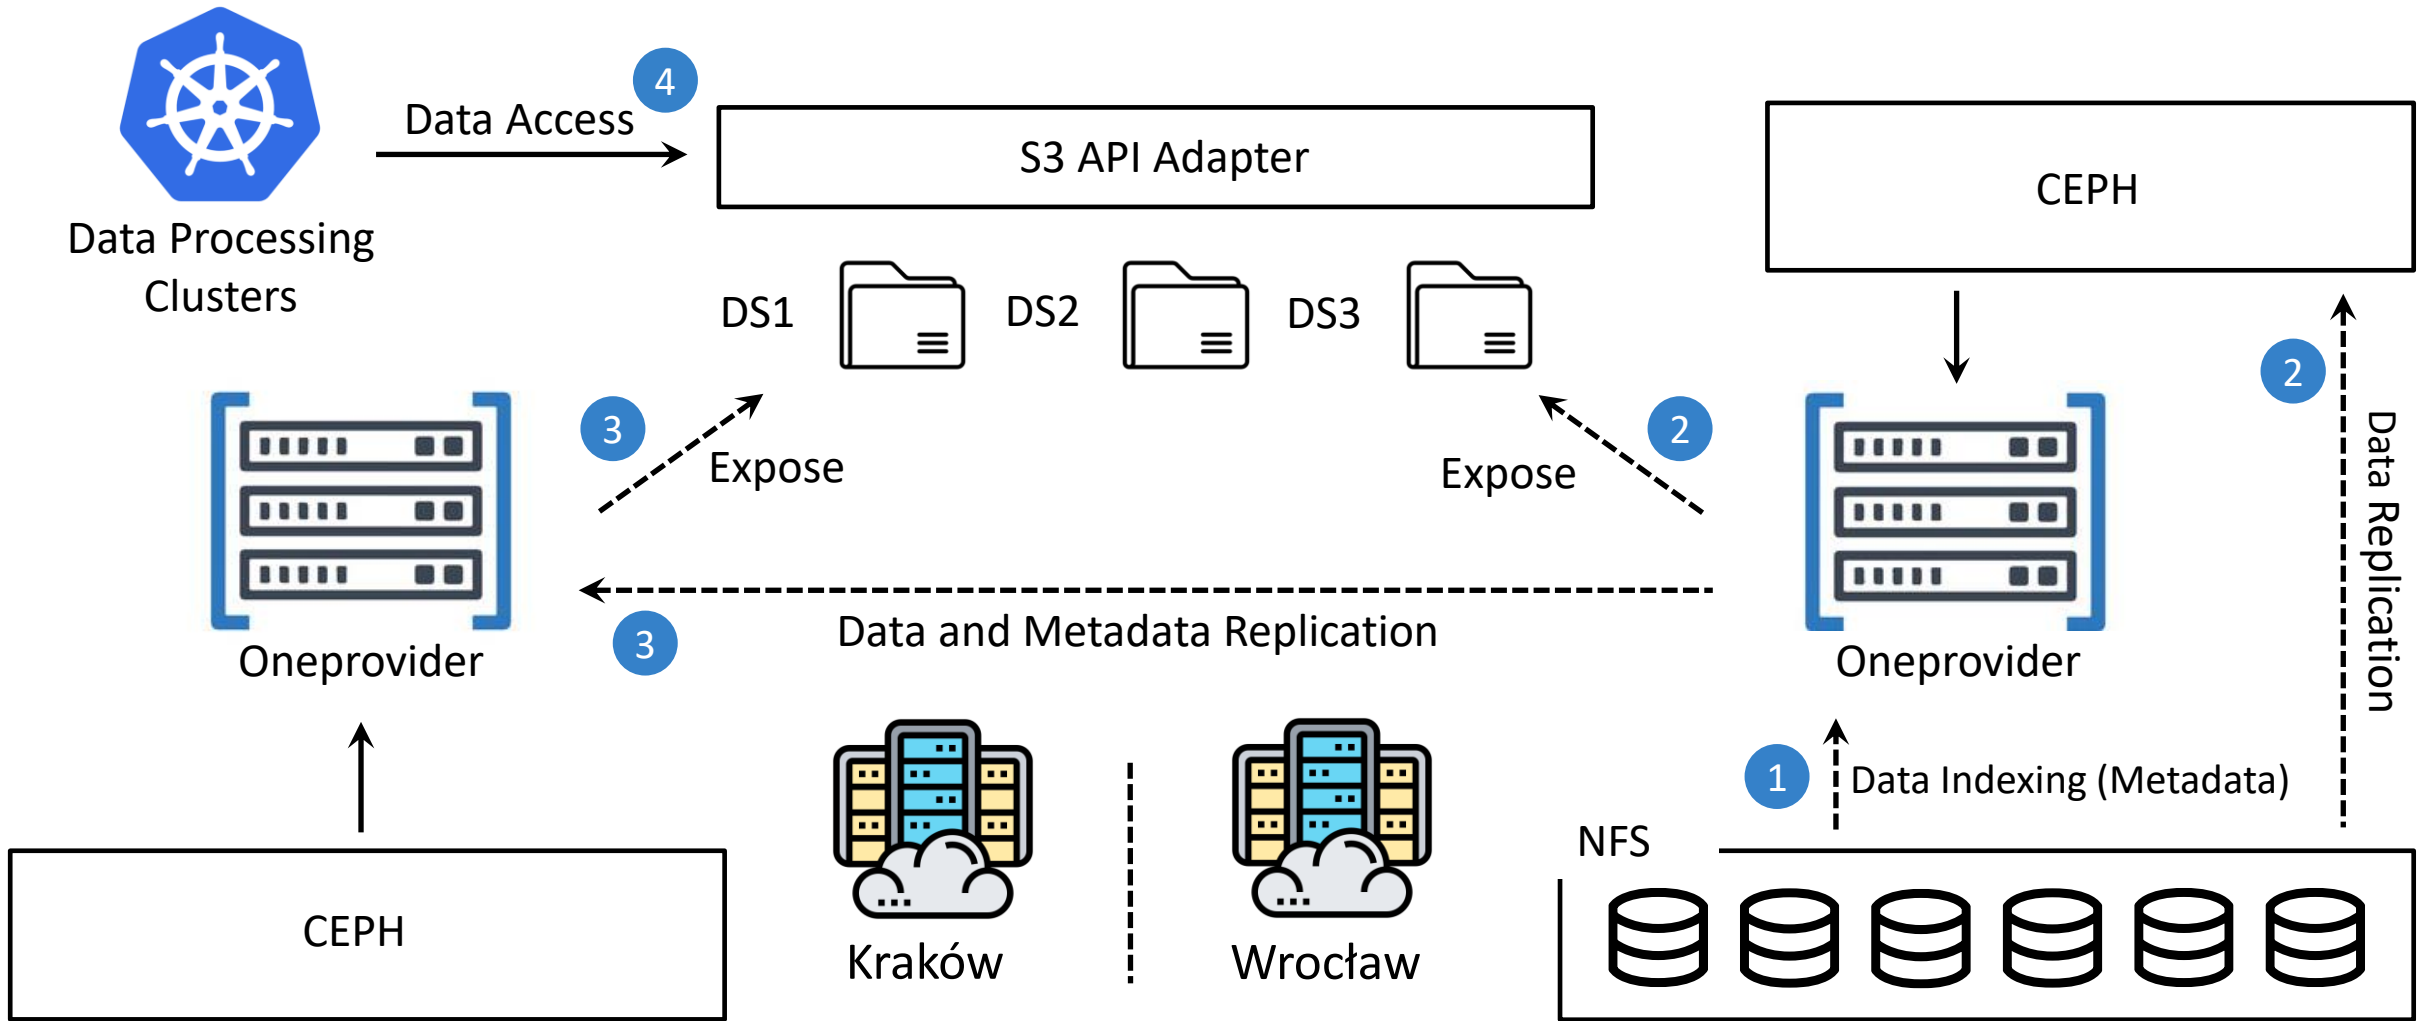
PROJECT REQUIRMENTS AND OBJECTIVE

Requirements:

- high availability of data (replicated across two sites)
- high availability of data access and processing services
- data processing (image compression, metadata extraction) to start as soon as possible
- data processing software requires S3 API for data access
- results available to public as soon as possible

Objective: provide a portal where general public can search and access all available data.





AUTOMATIC DATA INDEXING

The file indexing process does not copy any data, it simply creates the necessary metadata so that the files pre-existing on the storage are reflected and accessible in Onedata.

This functionality should be used in two setups:

- There is a legacy dataset located on the storage backend
- The data on storage backend is to be modified directly by third party applications, and the changes should be reflected in Onedata.

It is possible to configure the storage to:

- detect consecutive data changes
- manually triggering scans
- re-registering files

Set	A	B	C
Number of files	3934561	3180622	36790834
Number of directories	1	33667	154578
Total Size [TB]	2068	960	131
Time to Scan [h]	16.6	21.8	22.2

Auto storage import configuration


Auto import scan completed, next scan: 19:38:39

Start scan

IMPORT FILES PROCESSING


Minute Hour Day

Created Modified Deleted




FILE OPERATIONS QUEUE STATISTICS

Queue length



IMPORT THROUGHPUT

Throughput



DATA TRANSFERS

Replicate files on demand and on the fly.

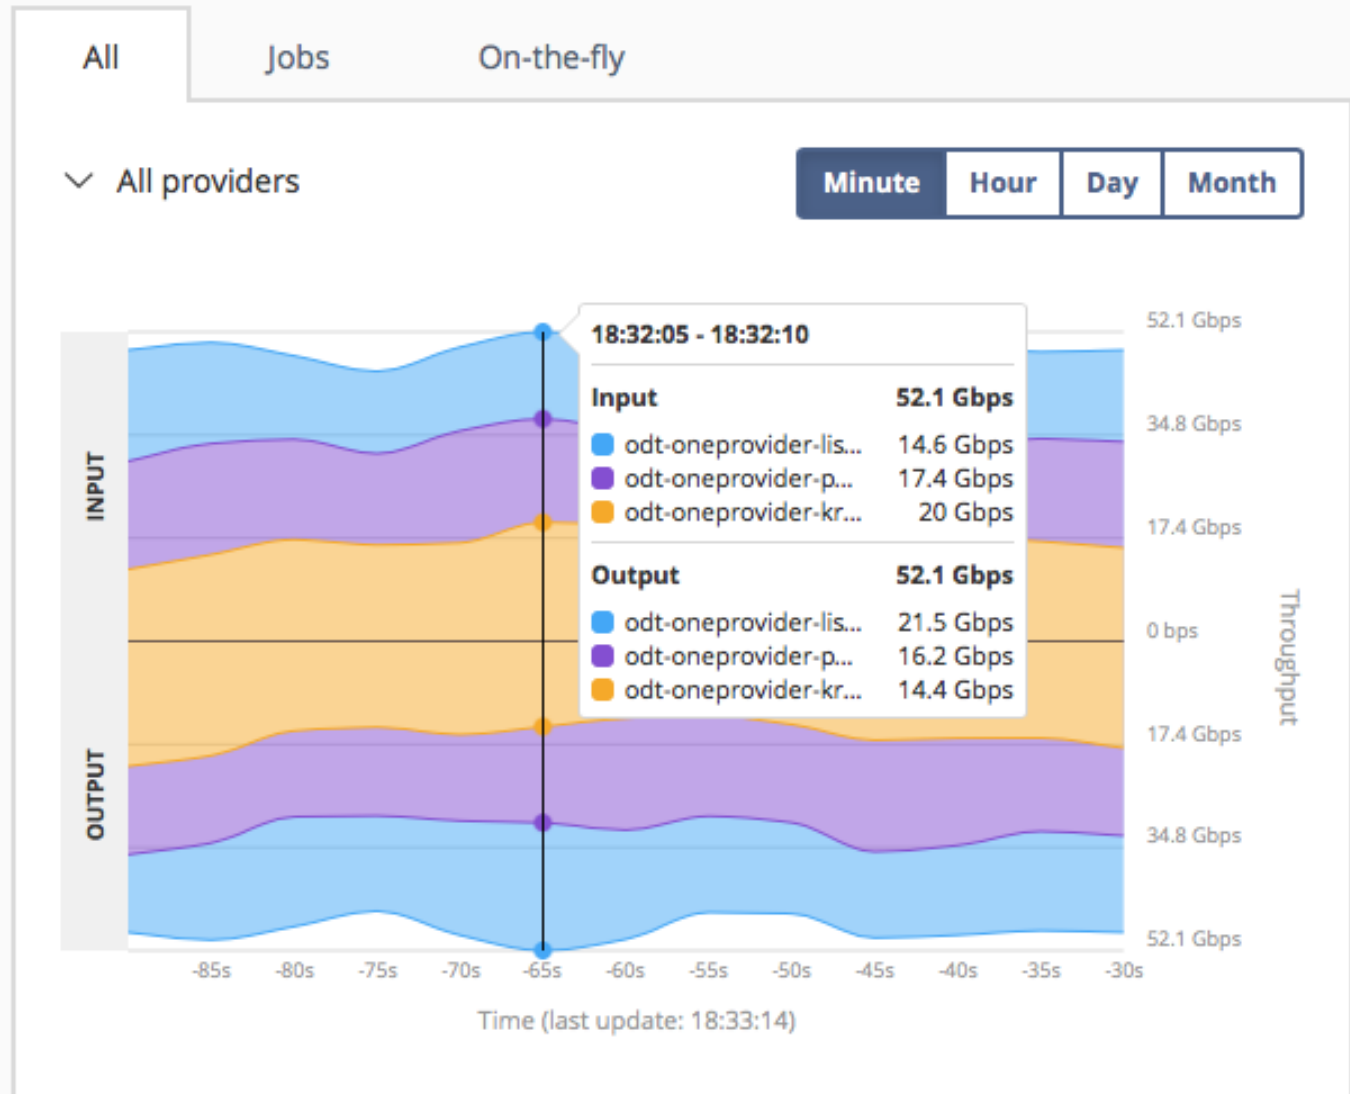
Migrate data between sites and storage backends on demand or with simple API interface.

Easily check location of your data using GUI or API.

Types of data transfers:

- replication — copying (only the missing) data to achieve a complete replica at the destination. The data is copied from one or more providers holding the missing blocks.
- eviction — removing replica(s) from the specified provider. This operation is safe and will succeed only if there exists at least one replica of every block on other supporting providers.
- migration — replication followed by eviction. Replicates the data to the destination provider and then evicts the replica from the source provider.

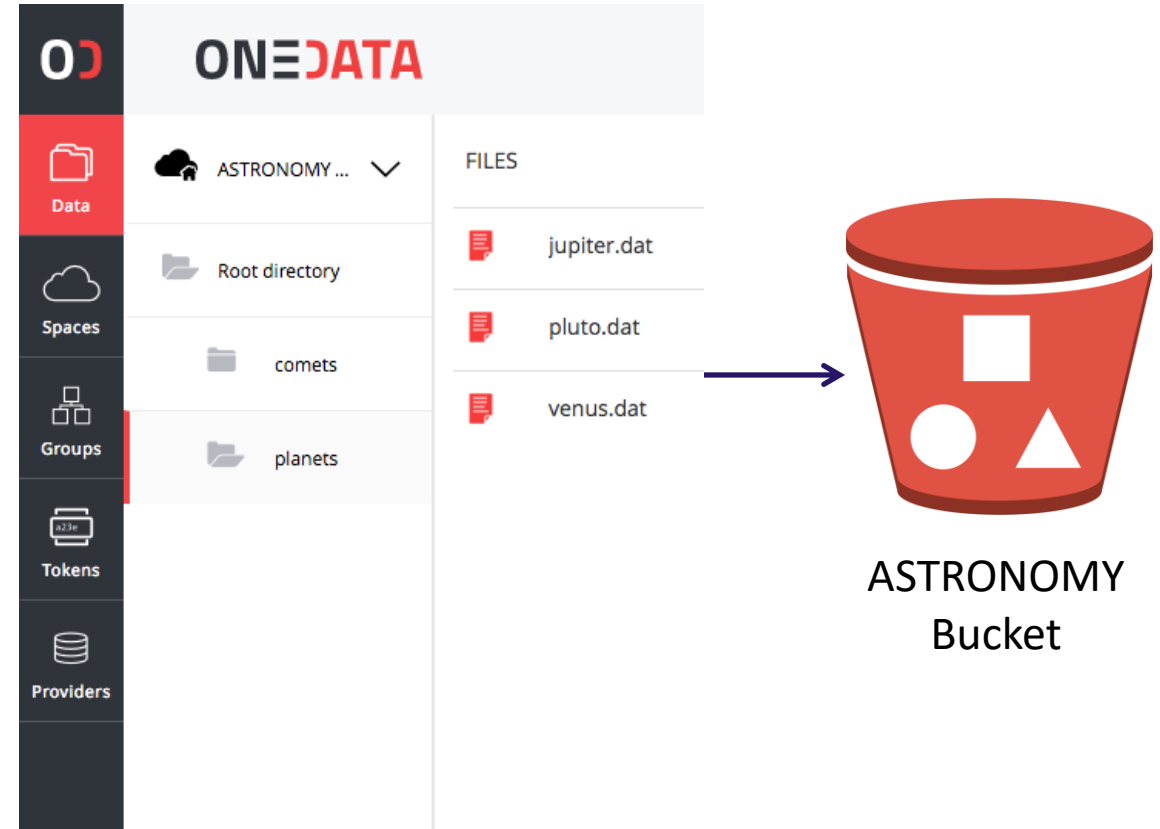
DATA TRANSFERS



ONE3S S3 API Adapter

OneS3 is a scalable S3 implementation:

- users can access their data spaces via AWS S3 interface,
- exposes storage resources managed by a specific Oneprovider,
- can be scaled to any number of nodes,
- it handles all data access to storage directly, while Oneprovider instances will only handle metadata management.



Portal z otwartymi zasobami kultury i nauki



lampa



Szukaj z warstwą tekstową



Szukaj tylko w tytułach





Zbiory Muzealne

Zespół "Stroma 5" (Żygulska, Żątowska, Karbowska) autor

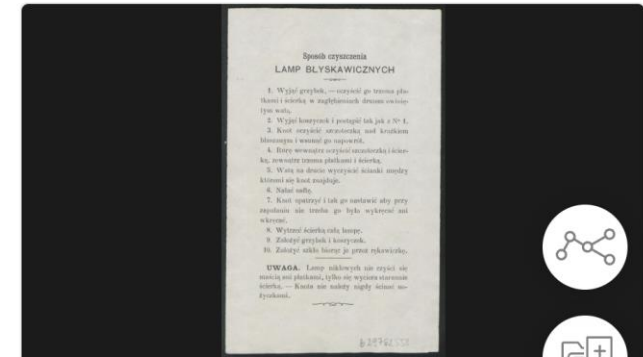
Lampa



Zbiory Muzealne

Zespół "Stroma 5" (Żygulska, Żątowska, Karbowska) autor

Lampa



Druki ulotne

Sposób czyszczenia lamp
zwyuczajnych.

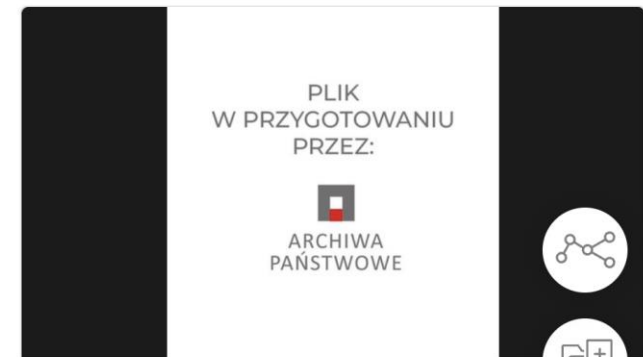
[193-]



Zbiory Muzealne



Zbiory Muzealne



T — dokumentacja techniczna

FOURTH USECASE:

**3D DIGITALIZATION AND PUBLICATION OF
CULTURAL HERITAGE OBJECTS**

PROJECT REQUIRMENTS AND OBJECTIVE

Requirements:

- university friendly authentication and authorization
- simple data upload interface
- annotating data with metadata
- data versioning and archiving
- integrated pipeline for processing 3D models
- data publication as open data and integration with Europeana

Objectives:

- provide a web based 3D viewer for showcasing 3D objects
- publication of 3D objects to Europeana

ONE
WORLD



EGI DATAHUB
ZONE

Sign in

with your identity provider



Onezone has been upgraded to version 21.02.1 on 18th May 2023.

Note for users: access to the PLAYGROUND space with a 30GB shared quota is given to all users via EGI Check-in.

Note for administrators: it is recommended to upgrade Oneprovider instances to version 21.02.1.

DataHub is an EGI service provided by CYFRONET, receiving funding from the EGI Foundation (EGI.eu) and the EGI-ACE project (Horizon 2020) under Grant number 101017567 | Powered by OneData.



Check-in

Choose your academic/social account

Search...

- 29 Mayis University
- A*STAR - Agency for Science, Technology and Research
- A. T. Still University
- AAF Virtual Home
- aai.lab.maeen.sa
- AAI@EduHr Single Sign-On Service
- Aalborg University
- Aalto University
- Aarhus School of Architecture
- Aarhus School of Marine and Technical Engineering
- Aarhus University
- AARNet
- Aba Teachers University
- Abertay University
- Aberystwyth University
- ABES - French Bibliographic Agency for Higher Education
- Abingdon and Witney College
- Absalon University College
- Academia Cotopaxi - cotopaxi
- ACADEMIA d.o.o.

or

Bitbucket D4SCIENCE esgi sso eduHr LOG-IN

B2ACCESS Facebook GitHub Google eduTEAMS

IGTF LinkedIn ORCID umbrella WeChat

ria

Can't find your identity provider?

INTEGRATED METADATA MANAGEMENT

Filesystem attributes — built-in filesystem metadata such as file size, creation and modification timestamps, POSIX permissions, etc. all files and directories can have a custom user metadata

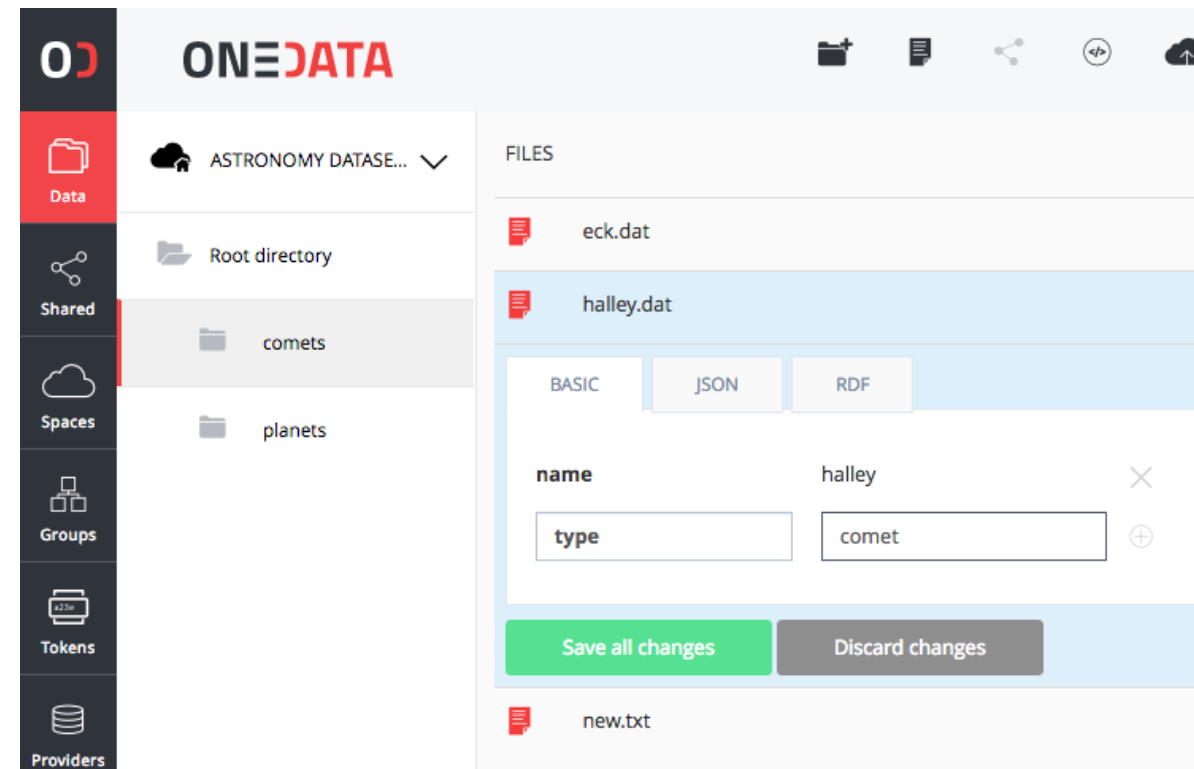
Custom metadata — arbitrary, user-defined information:

- extended attributes — a.k.a. basic attributes or xattrs; key-value pairs, compatible with POSIX extended attributes,
- JSON document,
- RDF document (XML).

All types of custom metadata can **coexist** for the same file.

API for metadata management.

Data discovery based on metadata.





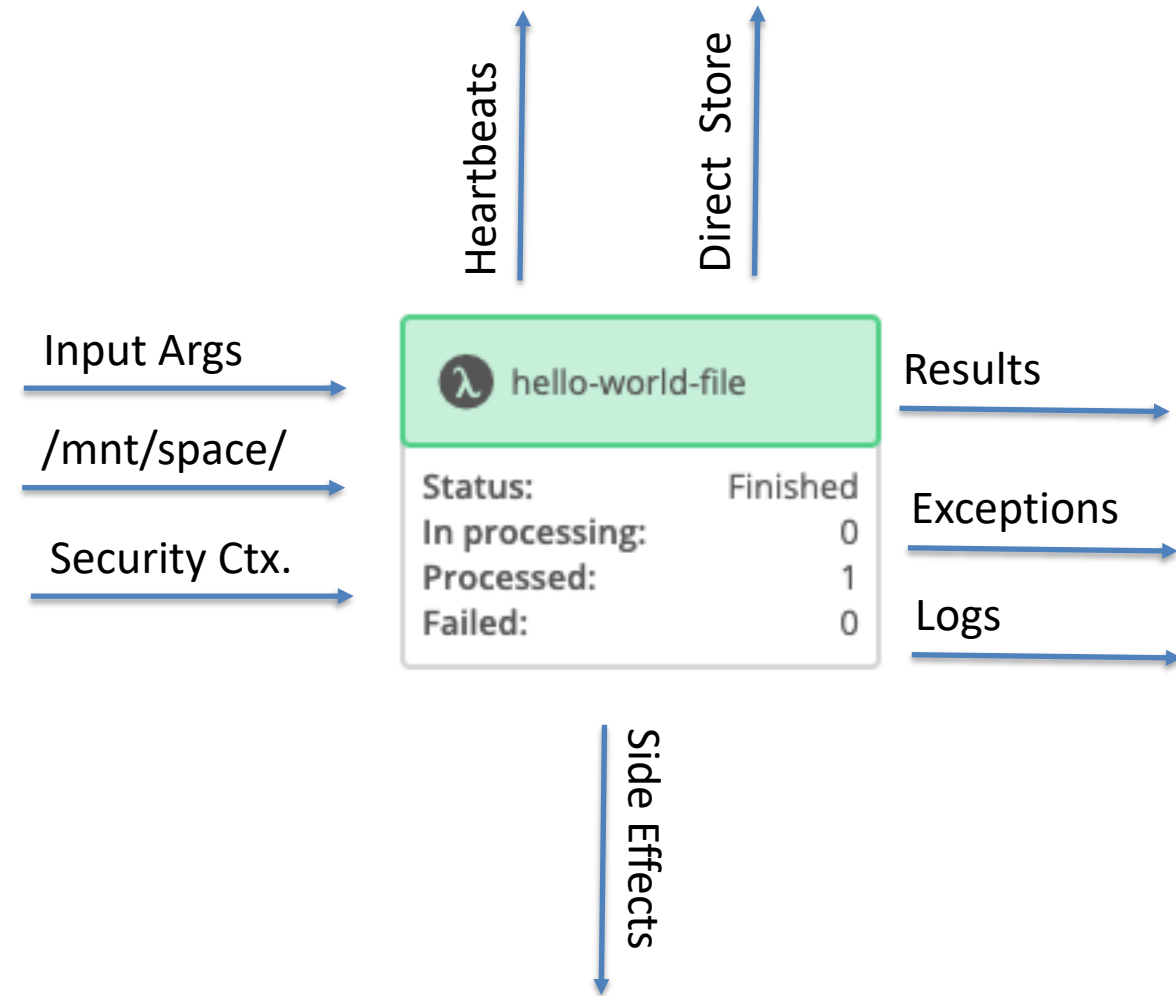
AUTOMATION ENGINE

FAAS LAMBDA

Lambda Anatomy

- **Input Arguments.** <Map>
- **Mount Space as File system.** <Oneclient> optional
- **Output Results.** <Map>
- **Exceptions.** <Map>
- **Logs.** <Map>
- **Side-effects.** e.g. REST-API calls
- **Heartbeats.** For long running lambdas
- **Stores Updates.** Direct operations on stores

- **Batch Mode.** Can work with batches of input arguments to speed up the process

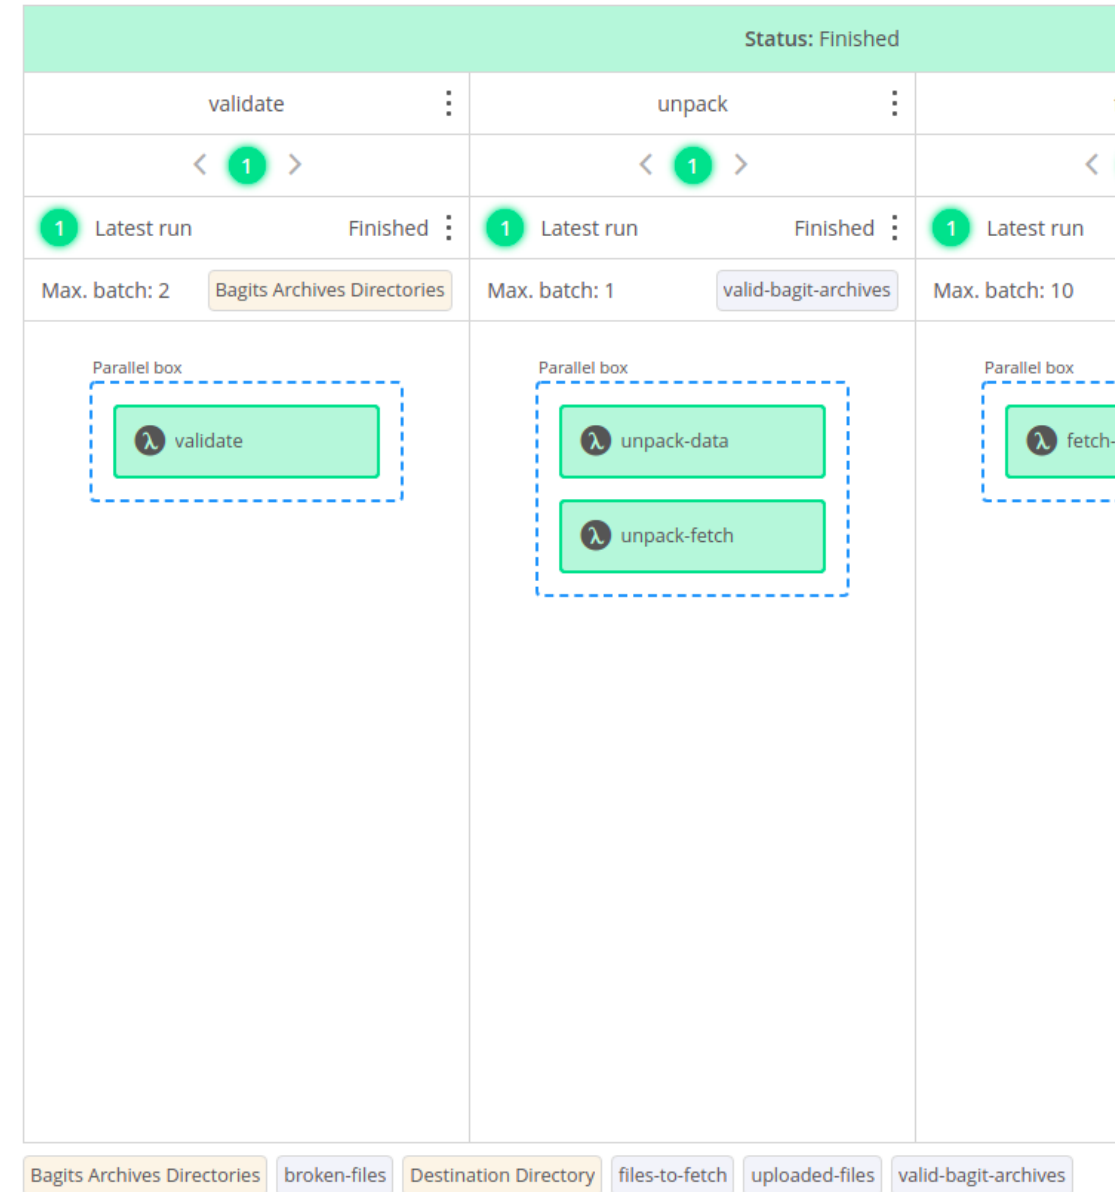


WORKFLOW

Workflow Anatomy

- **Lanes.** Iterates over Store and execute parallel boxes
- **Stores.** Input to to the workflow or produced during the workflow
- **Parallel Boxes.** Contains Lambdas which can be executed in any order
- **Lambdas.** Function which is called by mapping arguments

- Can be exported to JSON and reused by someone else



INVENTORY

Inventory Anatomy

- **Workflows.** Keep the list of workflows to be available for system users
- **Lambdas.** Keep the list of registered Lambdas
- **Members.** Access control
- **Import/Export.** Import full workflows into Inventory from JSON file

The screenshot displays the 'AUTOMATION' section of the Inventory interface. A vertical sidebar on the left contains navigation icons, with the 'Automation' icon highlighted in red. The main content area shows a list of automation workflows:

- Lukasz Inventory
- System Inventory** (highlighted in red)
- Workflows
- Lambdas
- Members

Below the list, the 'WORKFLOWS' section is visible, showing a search bar and three workflow cards:

- Bagit Extractor**: Automation workflow processing. Includes a table with 1 revision in 'Stable' state.
- Check Format**: Automation workflow for checking. Includes a table with 2 revisions: 1 'Draft' and 1 'Stable'.
- Checksum Calculator via POSIX**: Calculate checksums of file using. Includes a table with 1 revision in 'Stable' state.
- Hello World File**: The simplest possible function wh. Includes a table with 1 revision in 'Stable' state.

Rev.	State	Description
1	Stable	Stable Bagit Extra

Rev.	State	Description
2	Draft	Detecting new file
1	Stable	First version

Rev.	State	Description
1	Stable	Added MD5, SHA2

Rev.	State	Description
1	Stable	First version

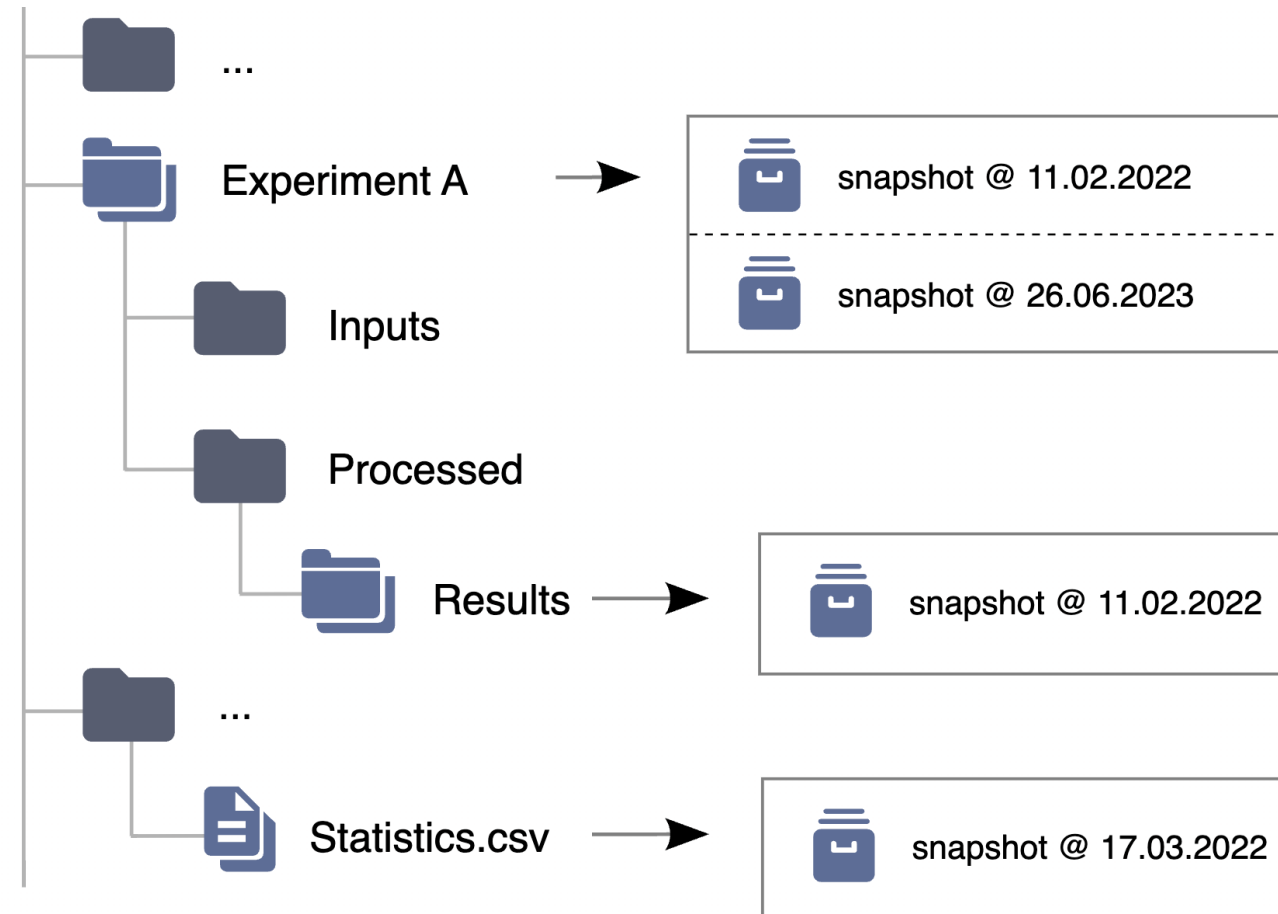
SNAPSHOTS – TOWARDS OPEN DATA

Datasets are essentially files or directories that have been marked by Onedata users as representing data **collections** relevant to them.

Datasets offer additional features, compared to regular files and directories:

- optional data and **metadata protection**,
- dataset structure tracking using the dataset browser,
- ability to create **persistent snapshots** of the physical dataset contents — archives.

Datasets can be nested to compose **hierarchical structures**

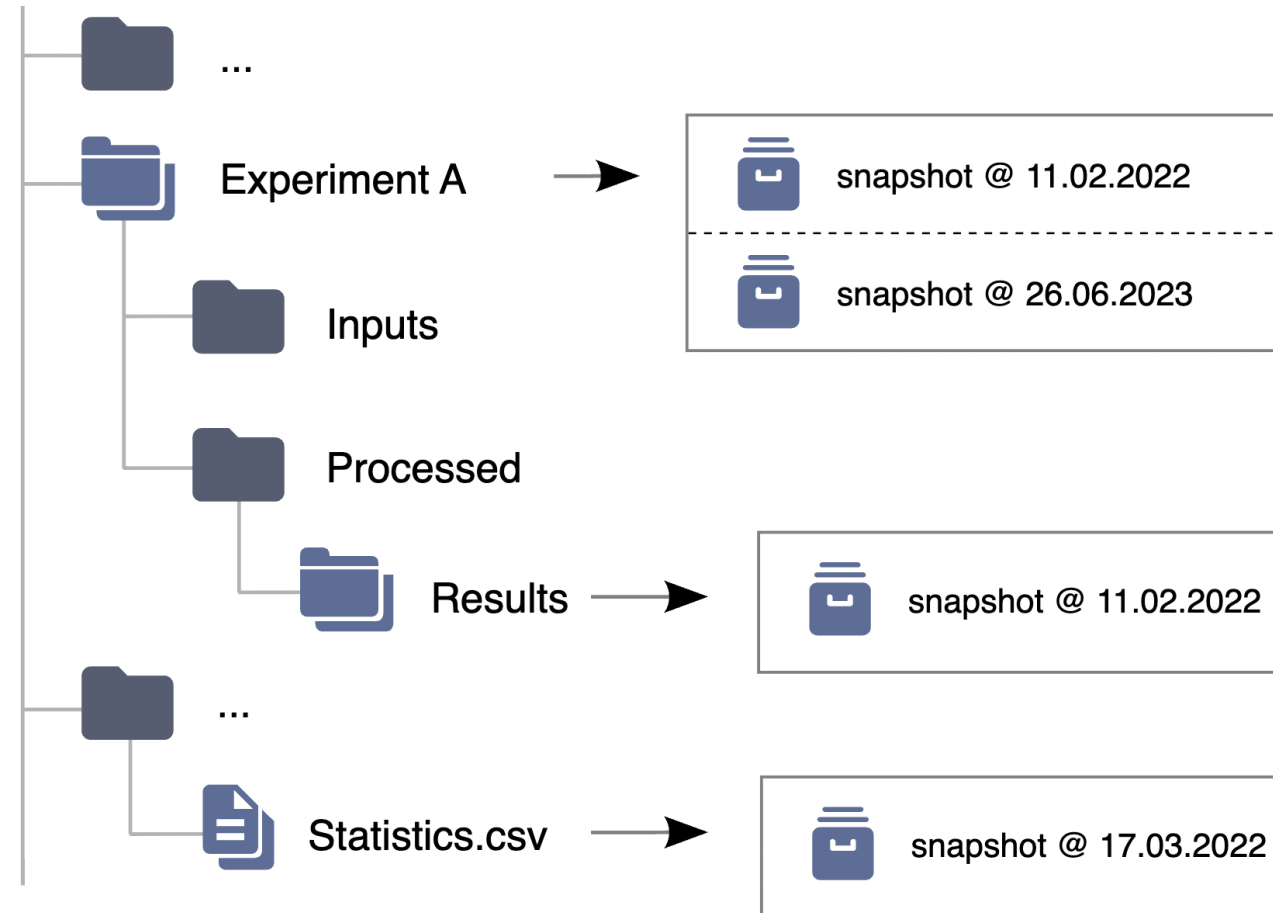


ARCHIVES – TOWARDS OPEN DATA

A snapshot of a dataset created at a certain point in time is called an archive.

The archive creation process comes with several options:

- different layouts — the structure of the files in the created archive; support for BagIt,
- creation of nested archives — hierarchically-created archives for datasets with nested structure,
- incremental archives — reusing the unchanged files between snapshots,
- possibility to follow symbolic links if they are present in the dataset



DATA SHARES – TOWARDS OPEN DATA

Files in Onedata can be made public and available for unauthenticated users.

A single file or directory can be shared using multiple links each can be removed at any time, giving the possibility to manage public access

The screenshot displays the ONEDATA web interface. On the left is a dark sidebar with navigation icons for Data, Shared, Spaces, Groups, Tokens, and Providers. The main content area is titled 'ONEDATA' and shows a 'SHARES' section with a folder icon and the name 'OxfordFlowerDataba...'. The right-hand panel is titled 'OxfordFlowerDatabase-FlowerSet1' and shows the following details:

- Path:** t7Y7mBQXgLv2RCRG_nvVF9dh268H86MickJ-0Hm3j84 > FlowerSet1
- Public URL:** https://datahub.egi.eu/share/h7tziSD2NmK2RSBkCaoA-_P2G7y... (with a copy icon)

Below the URL, there are tabs for 'BASIC', 'JSON', and 'RDF'. The 'BASIC' tab is active, showing a 'license' field set to 'CC-0' with a close icon (X). Below this is an 'Attribute' field with a 'Value' field and a plus icon (+). At the bottom of the settings panel are three buttons: 'Save all changes' (green), 'Discard changes' (grey), and 'Remove metadata' (red).

At the bottom of the interface, there is a section for 'FlowerSet1' with a 'FILES' sub-section containing a single file entry: 'image_0001.jpg' with a document icon.

PUBLICATION AS DUBLIN CORE

Dublin Core Metadata

Visual XML

Handle metadata is used to index the dataset in Open Data search engines and provide additional information for its consumers. ?

Title

+ Add another title

Ivory plaque with nativity scenes

Creator

+ Add another creator

The Hunt Museum as part of the Art of Reading in the Middle Ages project.

Description

+ Add another description

A 12th century Ivory plaque carved with scenes in three registers.

Top: the Annunciation (Mary and the Angel Gabriel) and the Visitation (Mary and Elizabeth). The Nativity with Mary lying on a bed and the Infant Christ just above her being warmed by the ox and ass. A figure stands on either side of this scene, possibly Salomé (the midwife) and Joseph.

Format

+ Add another format

glTF



Subject

+ Add another subject


Cultural Heritage & History

<http://vocab.getty.edu/aat/300011857>


PUBLICATION AS EUROPEANA DATA MODEL

 Provided Cultural Heritage Object 


Creator of the model Literal value Reference | Language:

 Value:


Description Literal value Reference | Language:

 Value:
Top: the Annunciation (Mary and the Angel Gabriel) and the Visitation (Mary and Elizabeth). The Nativity with Mary lying on a bed and the Infant Christ just above her being warmed by the ox and ass. A figure stands on either side of this scene, possibly Salomé (the midwife) and Joseph.

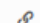
3D format Literal value Reference | Language:

 Value:


Subject Literal value Reference | Language:



 Value:

Subject Literal value Reference | Language:

 Reference:

Title Literal value Reference | Language:

 Value:

 Add property 



Lamp

Een grote gloeilamp en drie kleinere peertjes op een metalen plaatje. De grote lamp is bevestigd aan een verticale staaf met een stelschroef, om haar in hoogte te kunnen verstellen. Uit het glas van de grote lamp ontbreekt een stuk.

[Read more](#)

[Good to know](#) [All metadata](#)

Subject	03.06.0 Lamp ; Nernst
Type of item	Lamp
Date	1900 - 1915 ; 1900/1915
Providing institution	Rijksmuseum Boerhaave

EXTRA RESOURCES

Improved **documentation** (in making) <https://onedata.org/#/home/documentation>

Dedicated **demo mode** for easy sandbox deployment:

- <https://onedata.org/#/home/documentation/21.02/admin-guide/demo-mode.html>

Extensive **training materials** (4 day workshop!) covering majority of Onedata:

- <https://onedata.org/training>

user: training

password: Oneworkshop58

THANK YOU!

If any of this relates to your use cases, contact us!

Let's do some meaningful science together!

morzech@agh.edu.pl

info@onedata.org