Onedata

FaaS Data Processing with Onedata

Michał Orzechowski (CYFRONET AGH)

ACC Cyfronet AGH

WHO WE ARE?

- Group of developers bringing hybrid cloud open source platform to life
- 10+ years devoted development
- Our main goal is:
  - to deliver data management platform for large scale and distributed problems
  - to make the solution decentralized and eventually consistent in order build a mesh of data sources
  - to deliver virtual file system for hybrid cloud
- The work is supported by:

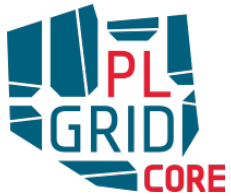

PROBLEMS ADDRESED BY ONEDATA

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 (DIO)
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

PROBLEM 1: MULTI-PROTOCOL TRANSPARENT ACCESS TO DATA IN MULTI-CLOUD ENVIRONMENTS

- Transparently access and create data in multi-cloud environments
- Care less about data locality, all your data are accessible wherever you go
- Use many protocols to access the same data

PROTOCOL HANDLERS (PLUGINS)

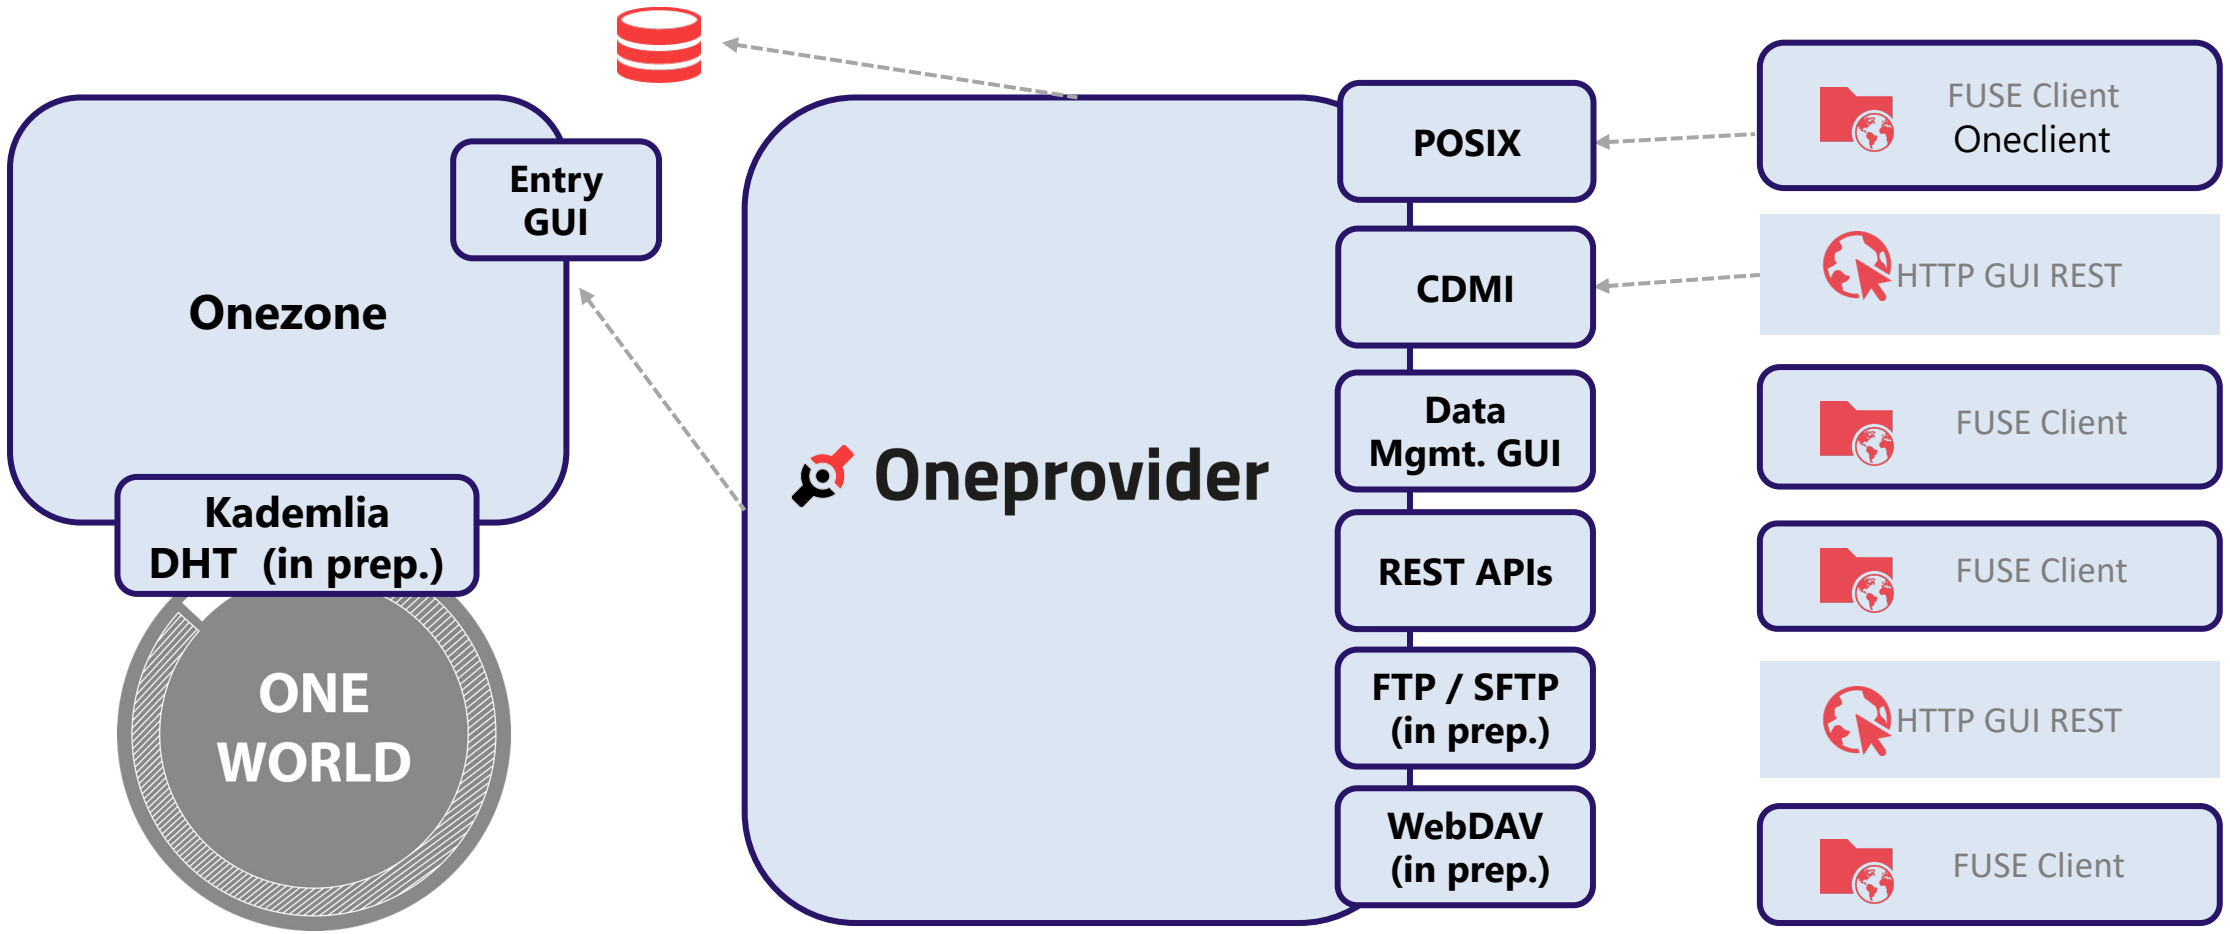

[...] BUT WE WANT POSIX

- 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

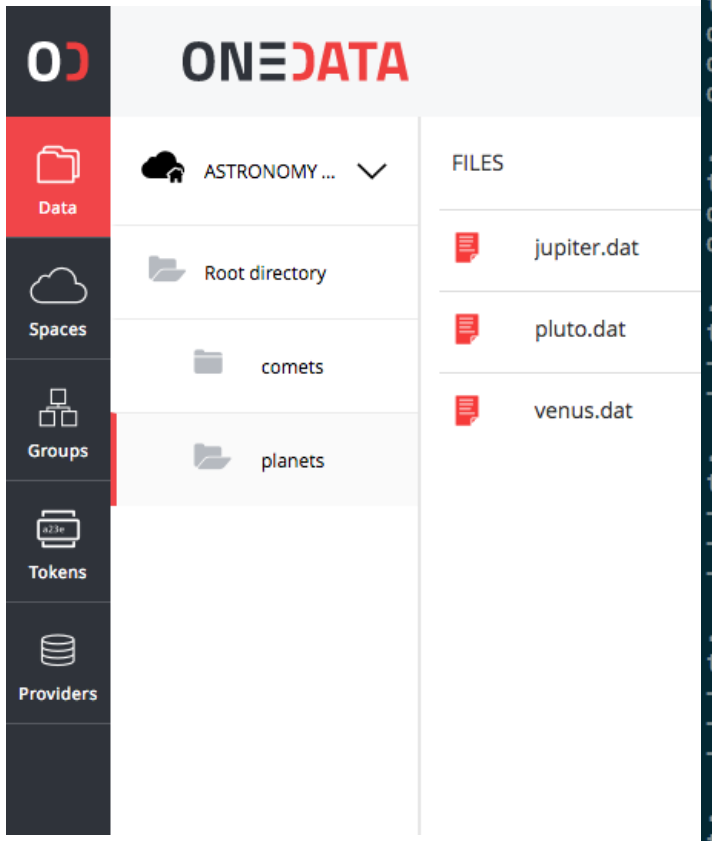


```
[root@1f87c053280e oneclient]# ls
Astronomy Datasets  Big Data Experiment  Cancer Data
[root@1f87c053280e oneclient]# ls -lR
.:
total 0
drwxrwx--- 1 root 1733762 0 Sep 26 19:19 Astronomy Datasets
drwxrwx--- 1 root 1337123 0 Sep 26 19:14 Big Data Experiment
drwxrwx--- 1 root  608582 0 Sep 26 19:18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

PROBLEM 2: HETEROGENEITY OF STORAGE TECHNOLOGIES

- Use the data protocols of your choice to access data wherever you go
- Minimize the problems of selection right storage technology to data centres operators
- Avoid cloud vendor locking

DIFFERENT TYPES OF STORAGES VIRTUALIZED

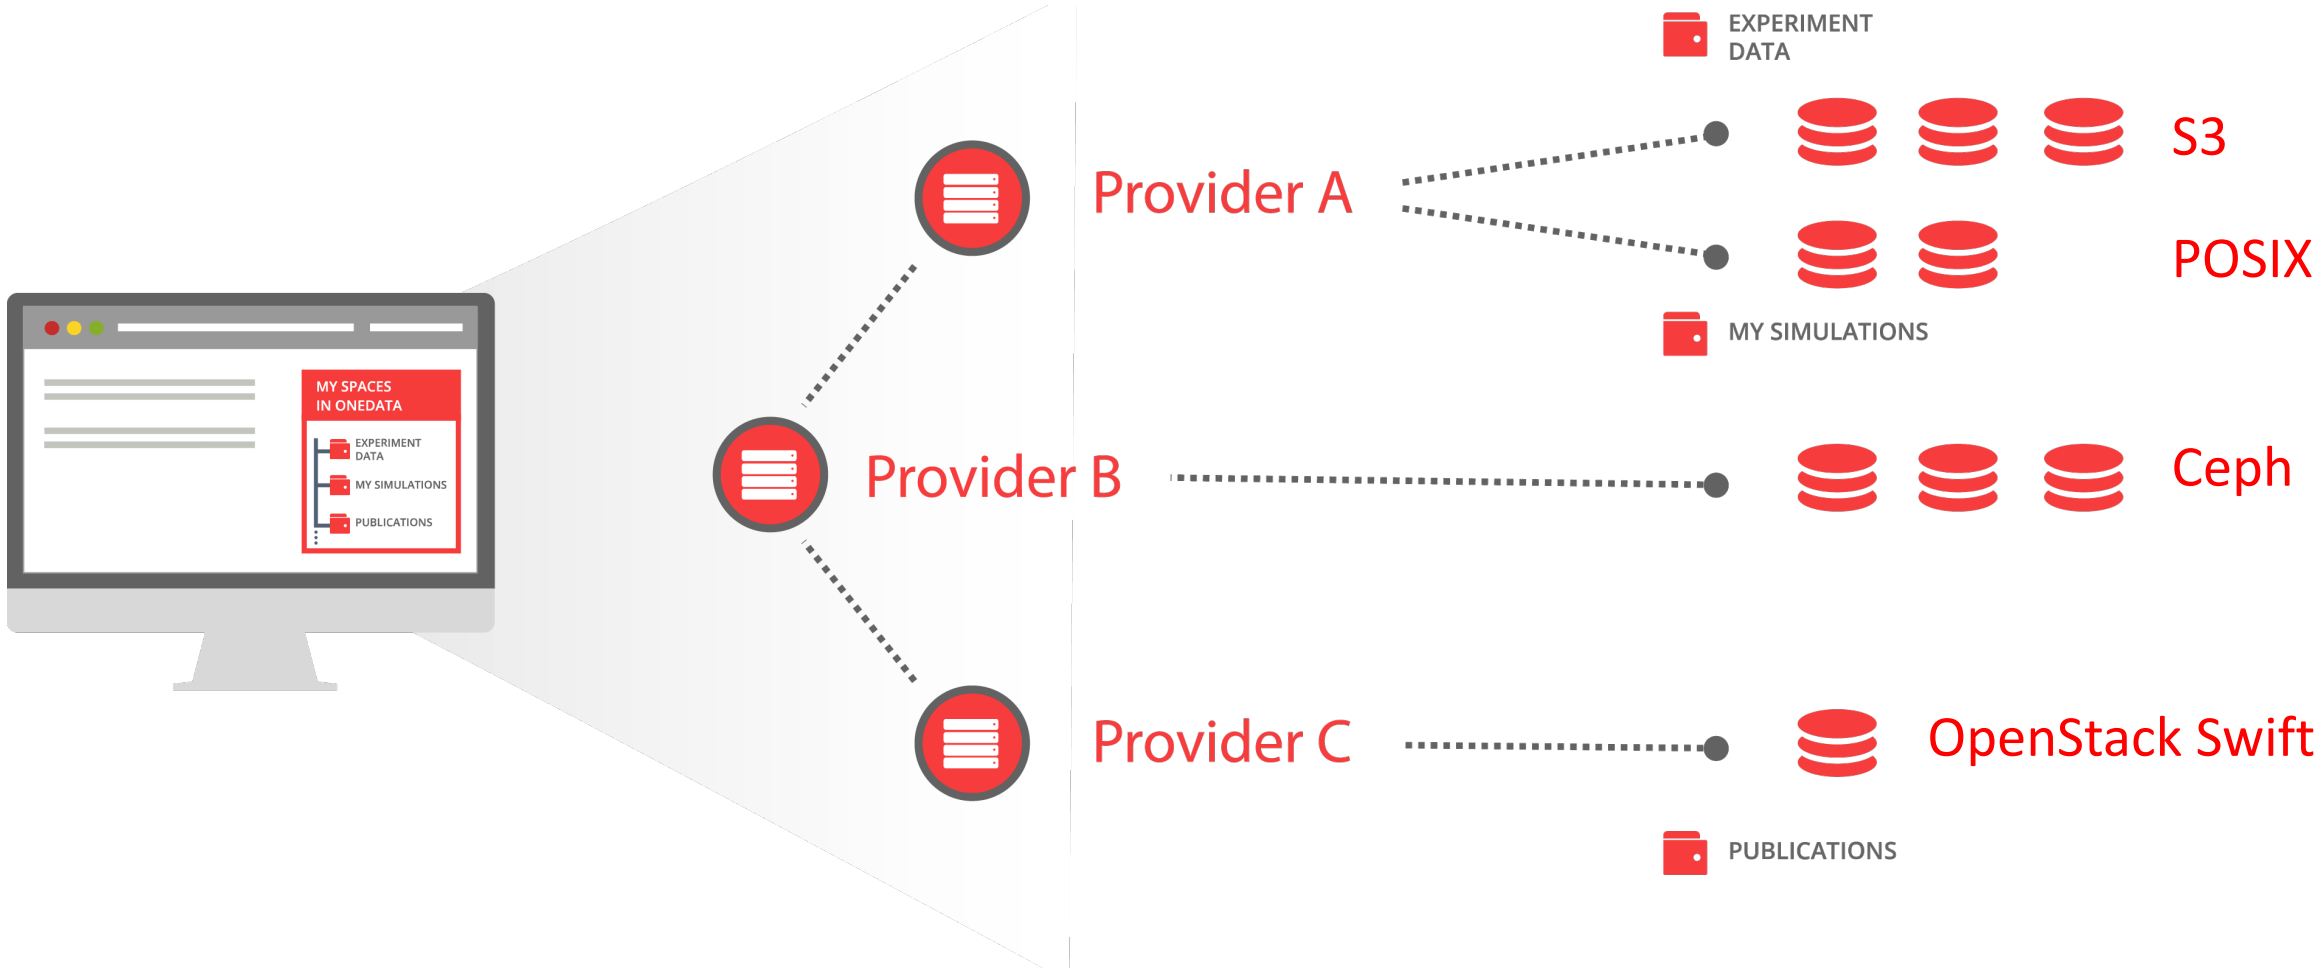

STORAGE SYSTEMS DRIVERS (PLUGINS)

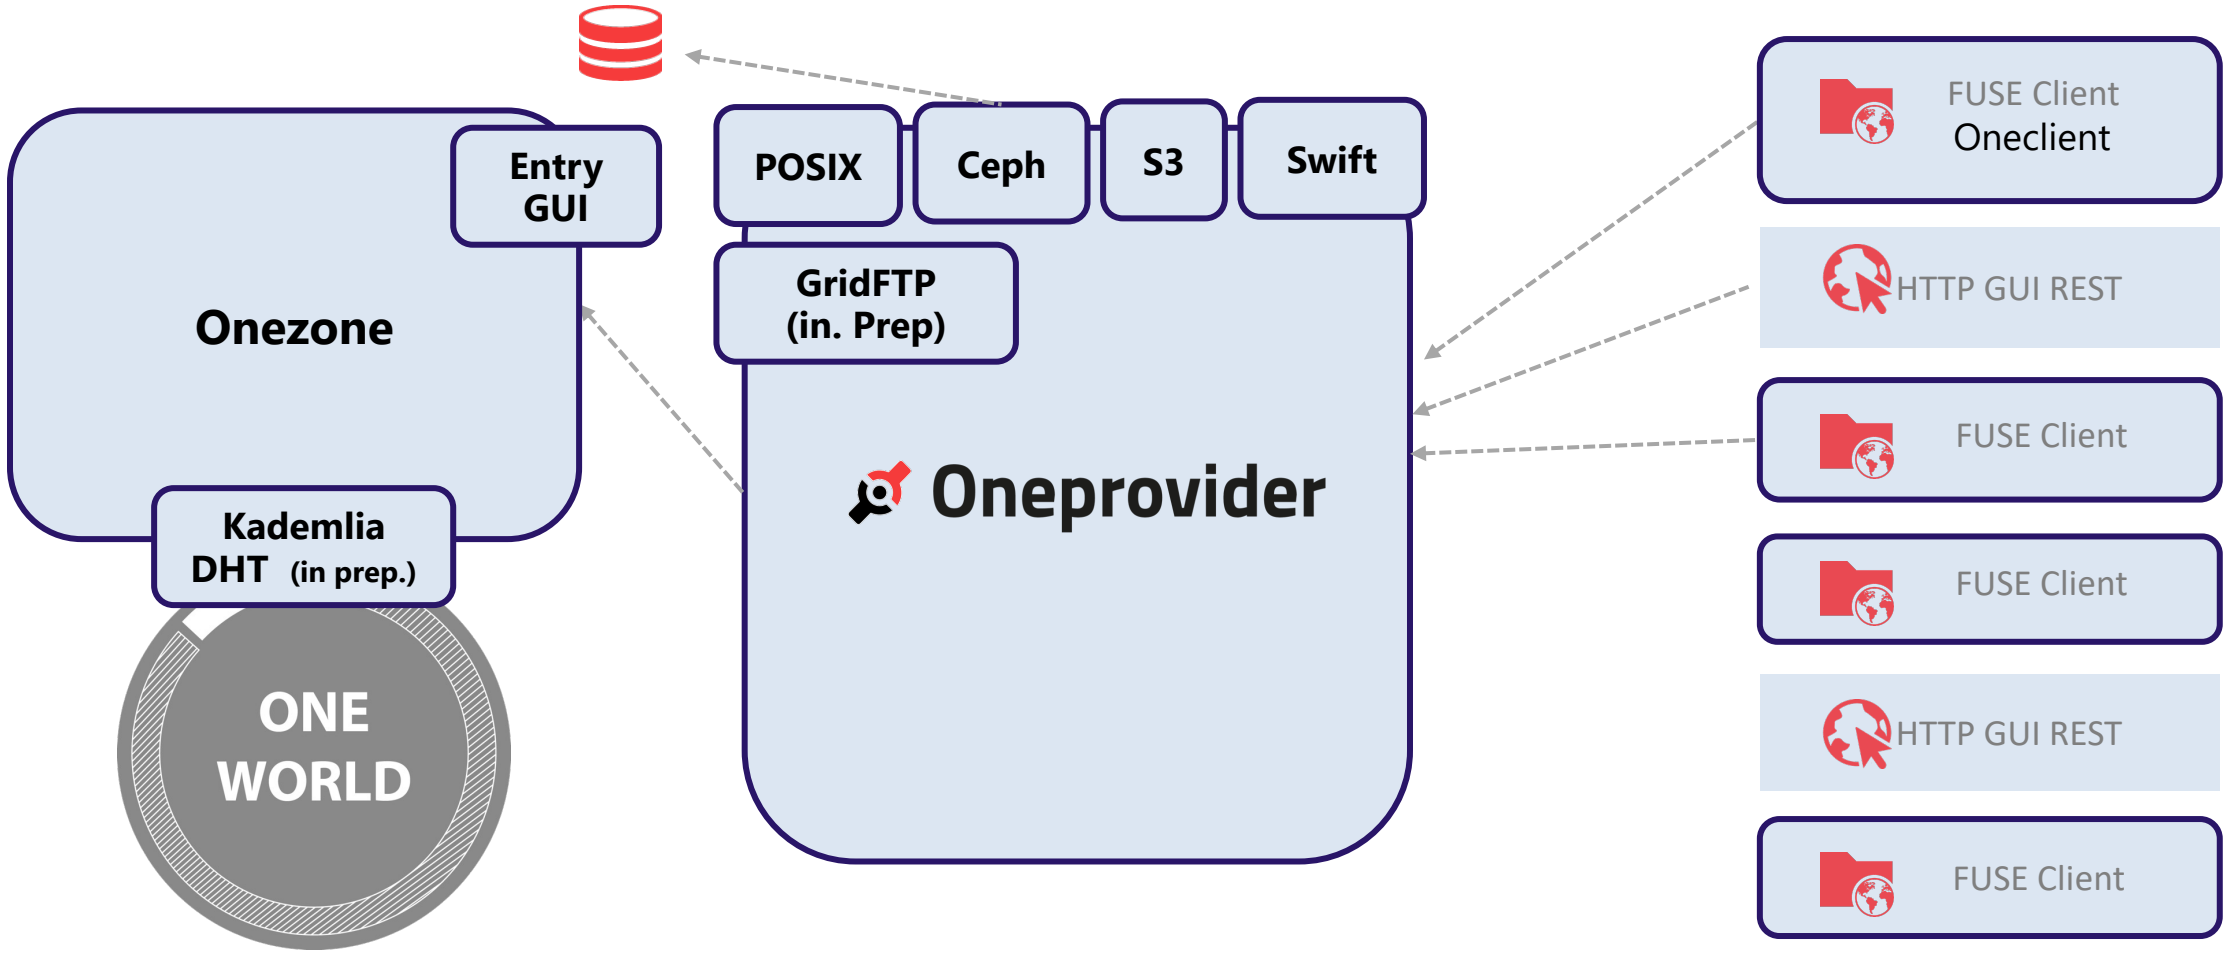

PROBLEM 3: REPLICA MANAGEMENT

- Replicate files on demand and on the fly without any additional effort
- Migrate data between sites on demand with simple API interface
- Easily check location of your data using GUI or API

REPLICAS MANAGEMENT SIMPLIFIED

- Manage files not Replicas
- File distribution between storage locations is underneath the file structure
- Replicas management on a chunk basis
- Missing chunks delivered on the fly
- API for replica management for pre-staging and implementing external data policy management

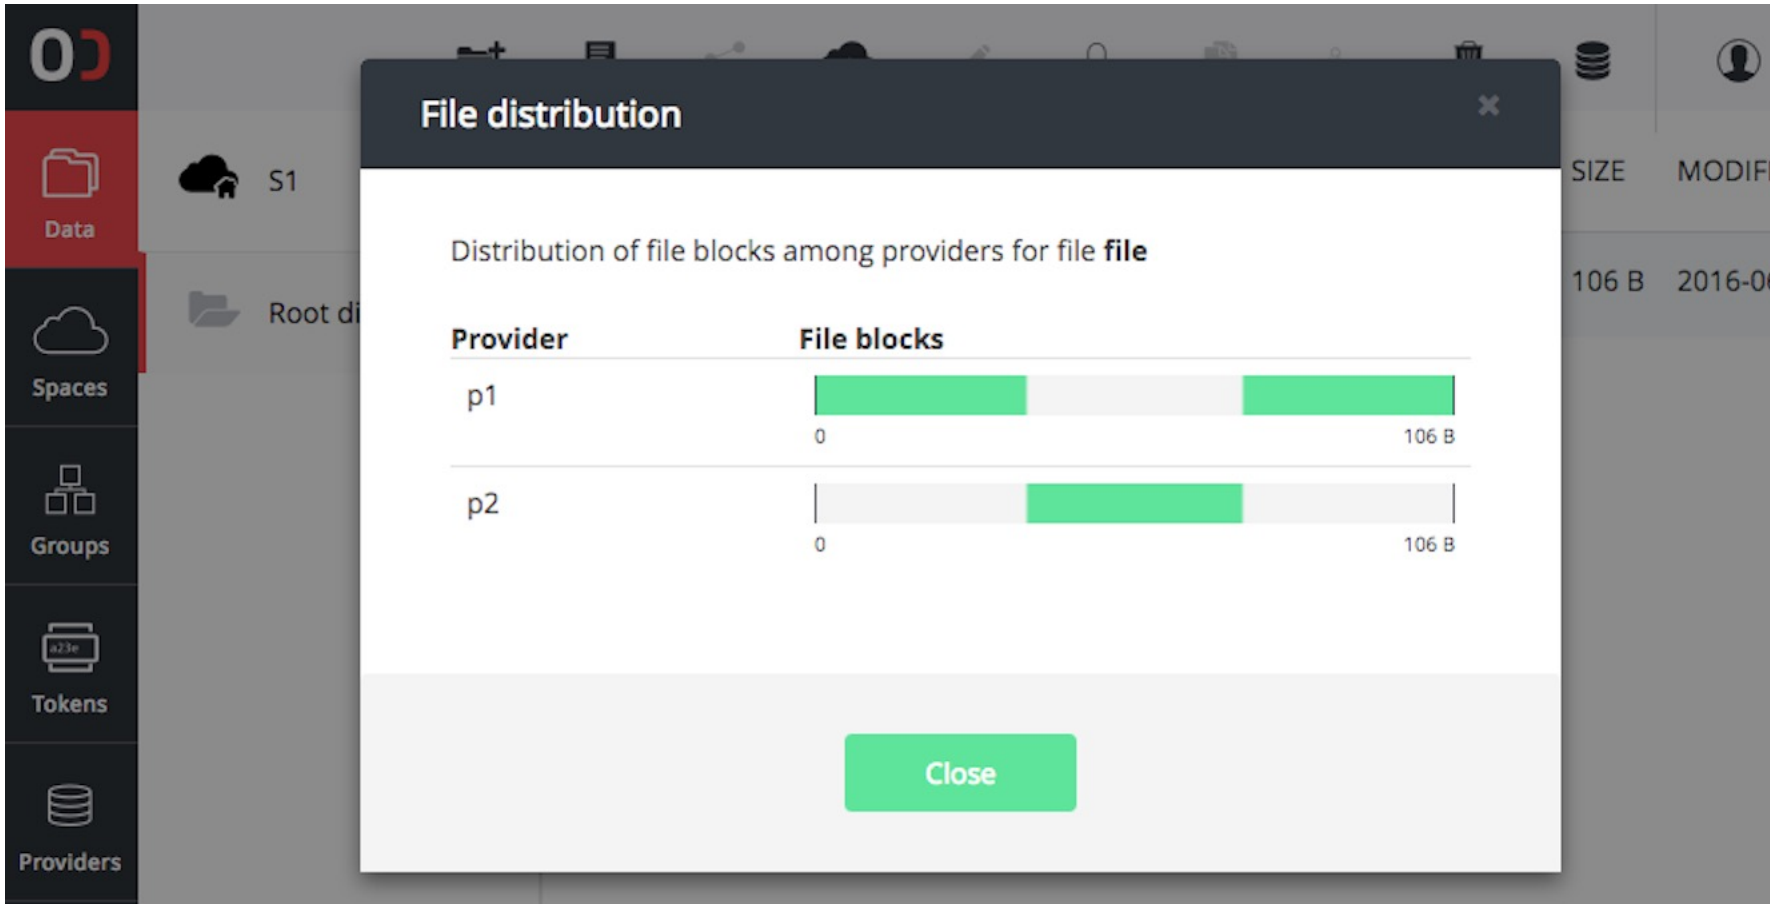

PROBLEM 4: EASY DATA SHARING WITHOUT BORDERS

- Share large scale data collections with other communities
- Enable your data to be shared in cross-federation scenarios
- Bring your data and tools as building blocks to European Open Science Cloud

EASY DATA SHARING

- Team-sharing
  - For groups
  - For individuals
  - Token based
- Cross-community data sharing
- Instant and ad-hoc data sharing
- *Thanks to effort supported by EGI Engage:*
  - *Open Data Publication*
  - *Handles (DOI) support*
  - *OAI-PMH*

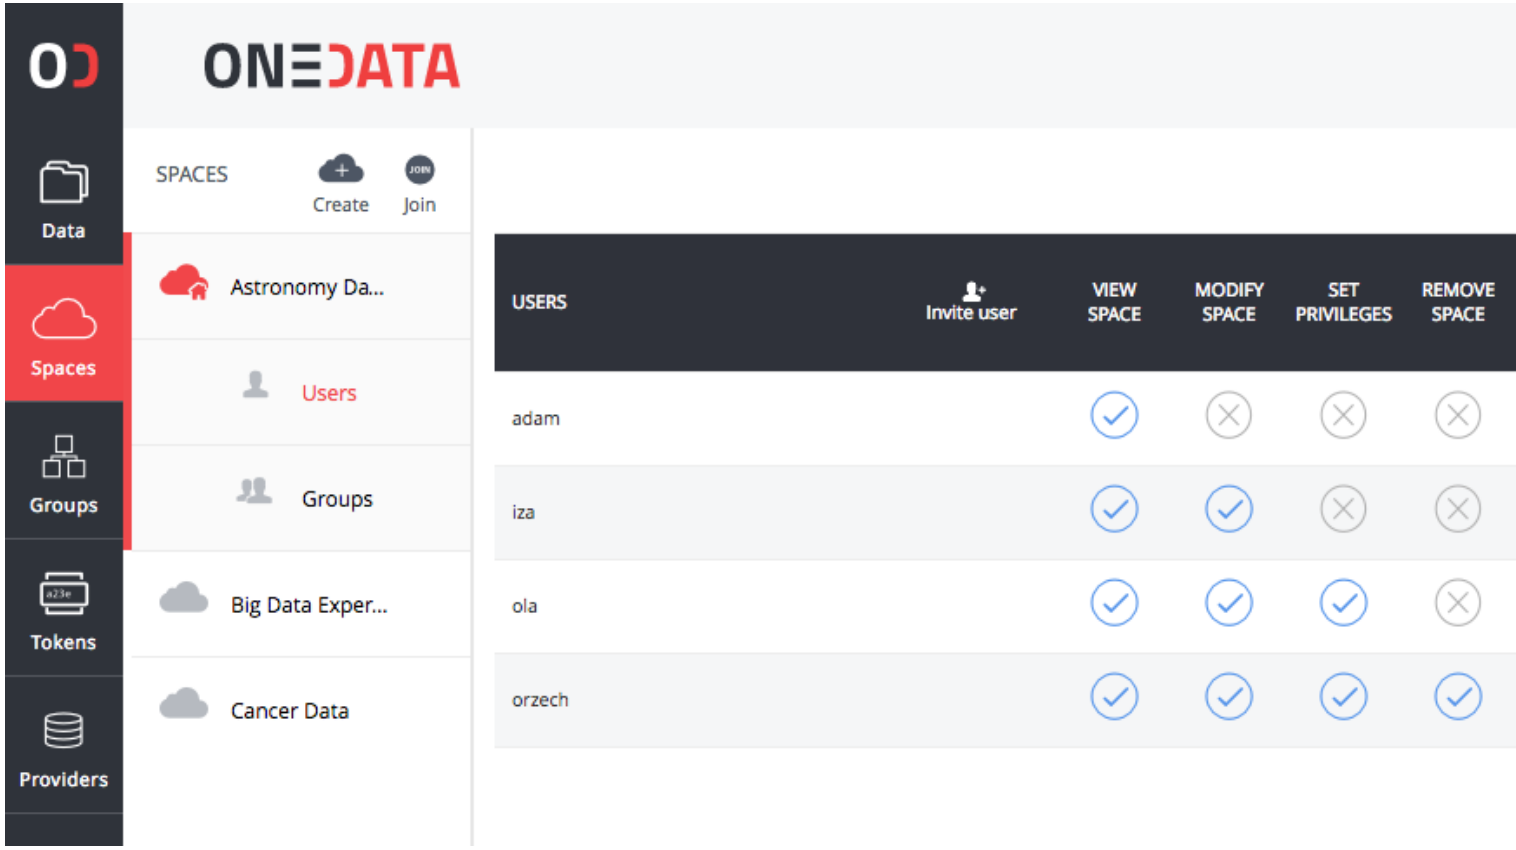

SHARES

OxfordFlowerDataba...

OxfordFlowerDatabase-FlowerSet1

Path t7Y7mBQXgLv2RCRG_nvvF9dh268H86MicKJ-0Hm3j84 > FlowerSet1

Public URL https://datahub.egi.eu/share/h7tziSD2NmK2RSBkCaoA-_P2G7y...

BASIC | JSON | RDF

license CC-0

Attribute | Value

Save all changes | Discard changes | Remove metadata

FlowerSet1

FILES	SIZE	MODIFICATION
image_0001.jpg	50.85 KB	2016-09-27 15:09
image_0002.jpg	41.24 KB	2016-09-27 15:09
image_0003.jpg	46.2 KB	2016-09-27 15:09
image_0004.jpg	30.16 KB	2016-09-27 15:09

PROBLEM 5: METADATA MANAGEMENT INTEGRATED WITH DATA MANAGEMENT PLATFORM

- Work with data and metadata in one system – avoiding problems of consistency
- Monitor metadata data changes trough API in order to feed external custom systems
- Advanced data discovery capabilities based on metadata

INTEGRATED METADATA MANAGMENT

- All files and directories can have a custom user metadata
- API for metadata management
- API for data discovery based on metadata
- Virtual Folders based on metadata tags
- Metadata formats:
  key-value, JSON, RDF

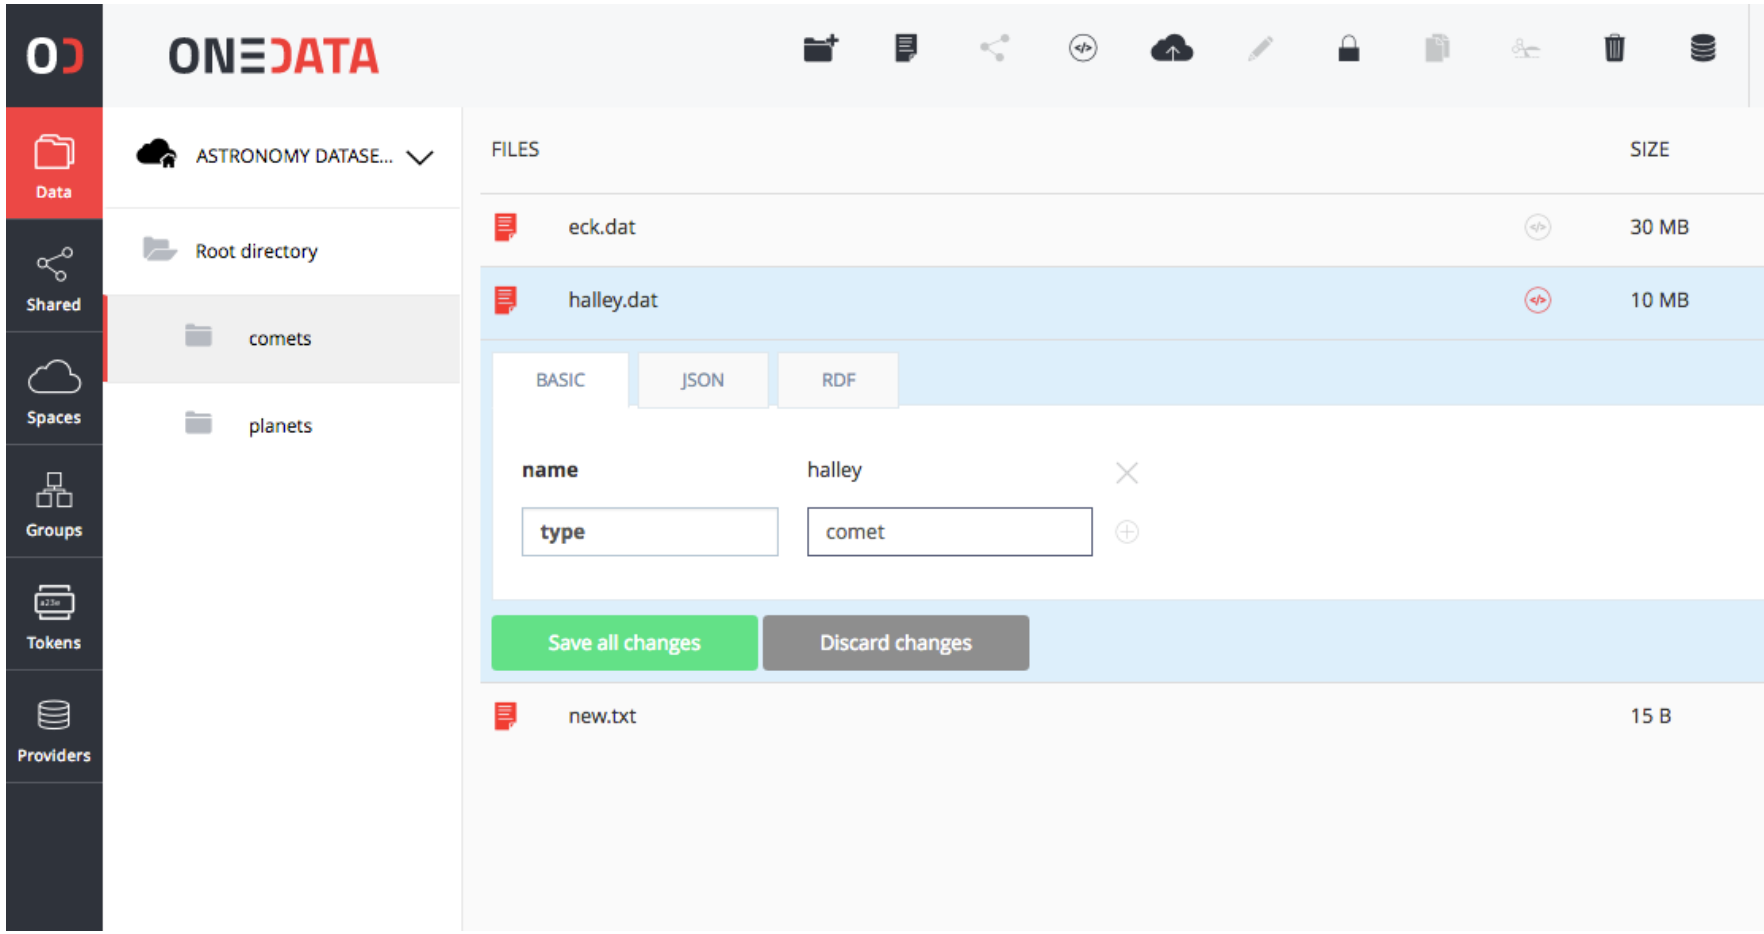

PROBLEM 7: EASY INTEGRATION USING API WITH EXTERNAL TOOLS

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

RICH COLLECTION OF APIs

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

PROBLEM 8: HIGH-THROUGHPUT PROCESSING

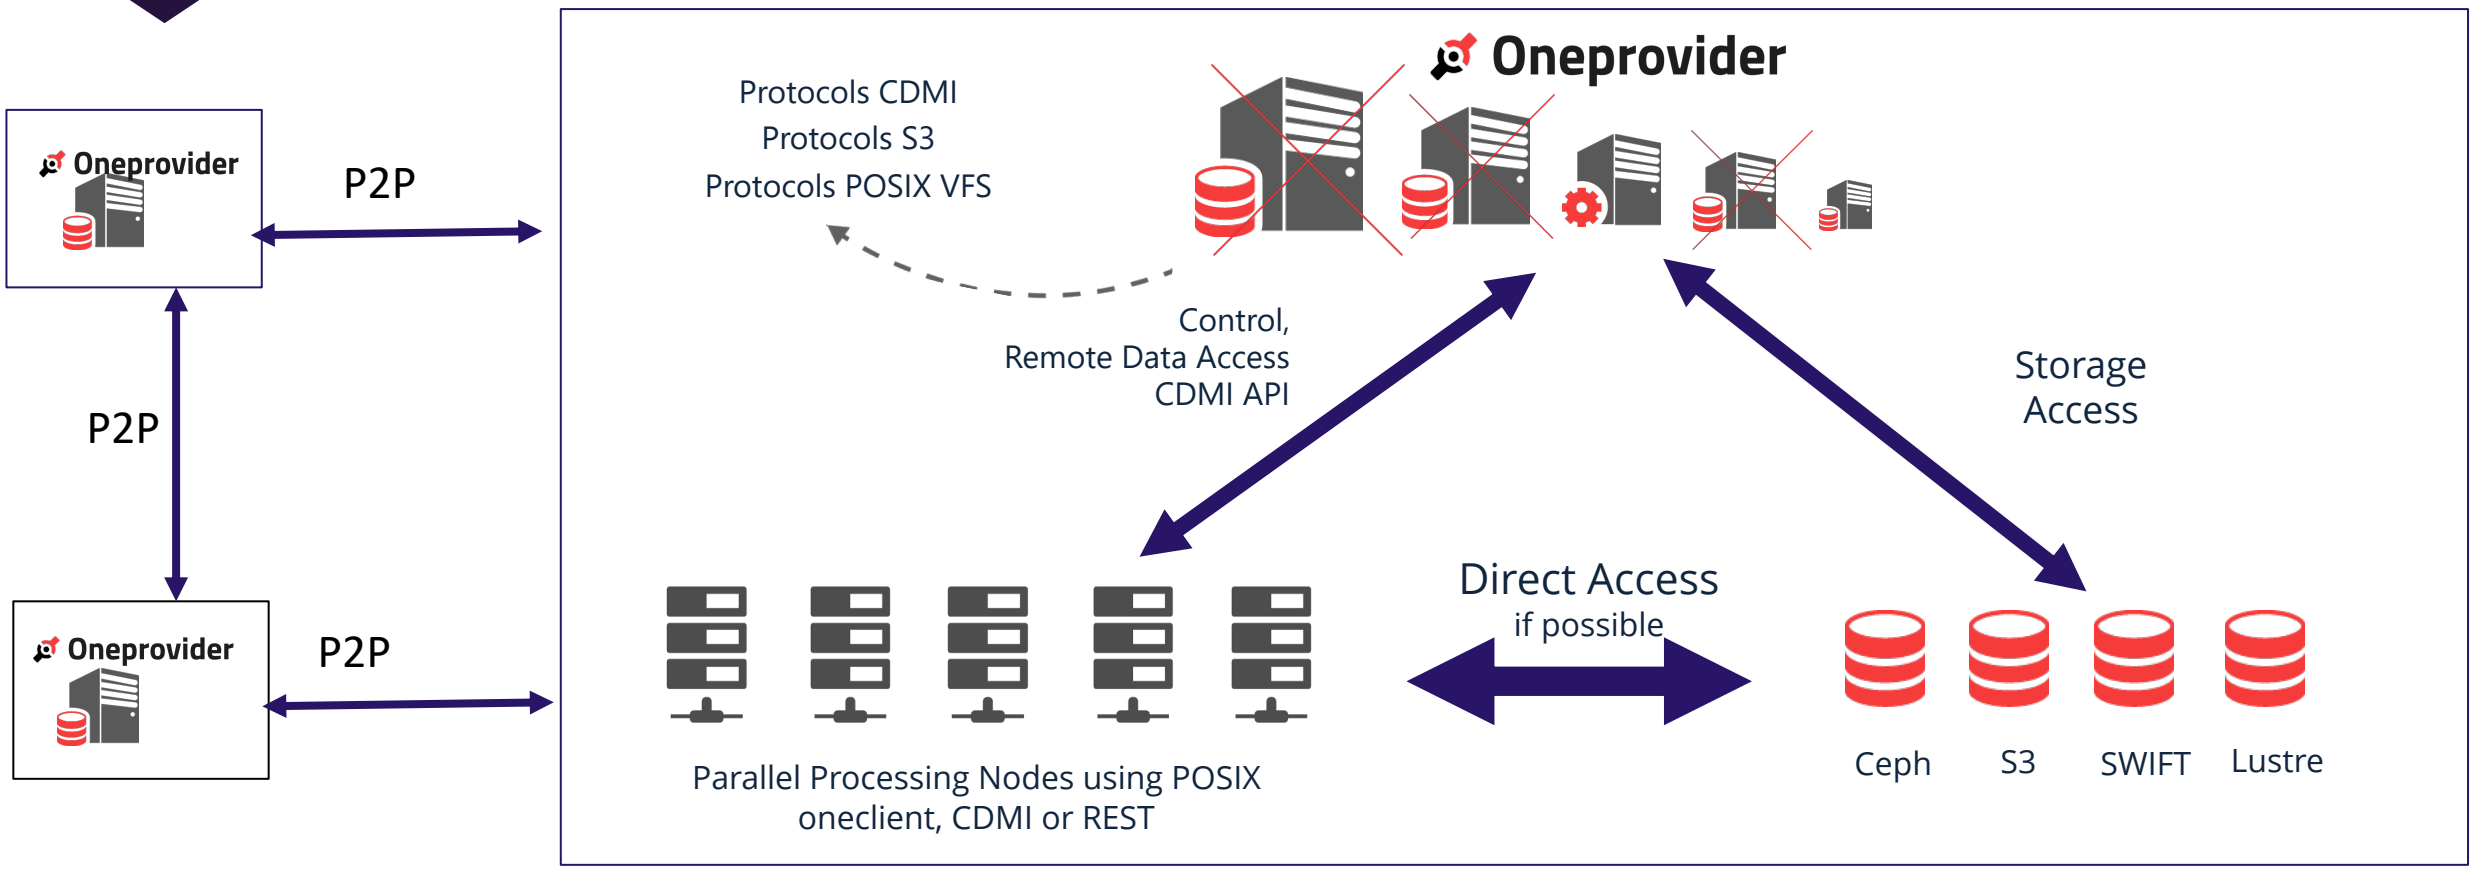

HIGH-THROUGHPUT TRANSFERS

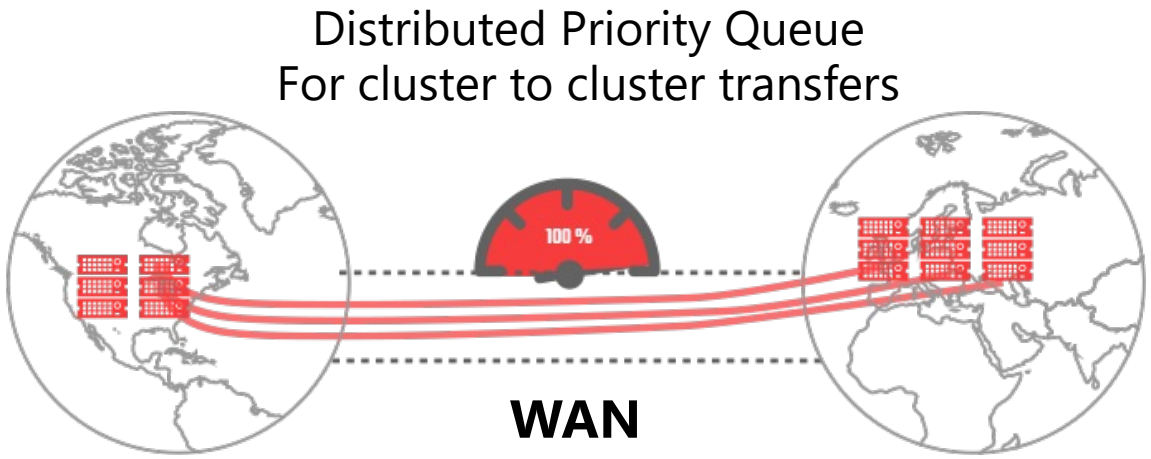

Transfer started by:

- User in GUI
- API-s
- Policy
- Access to Rmt. Data

Block-based transfer:

- Remote Data Access on the fly
- Pre-staging
- Data Migration
- Data Replication

New Features

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

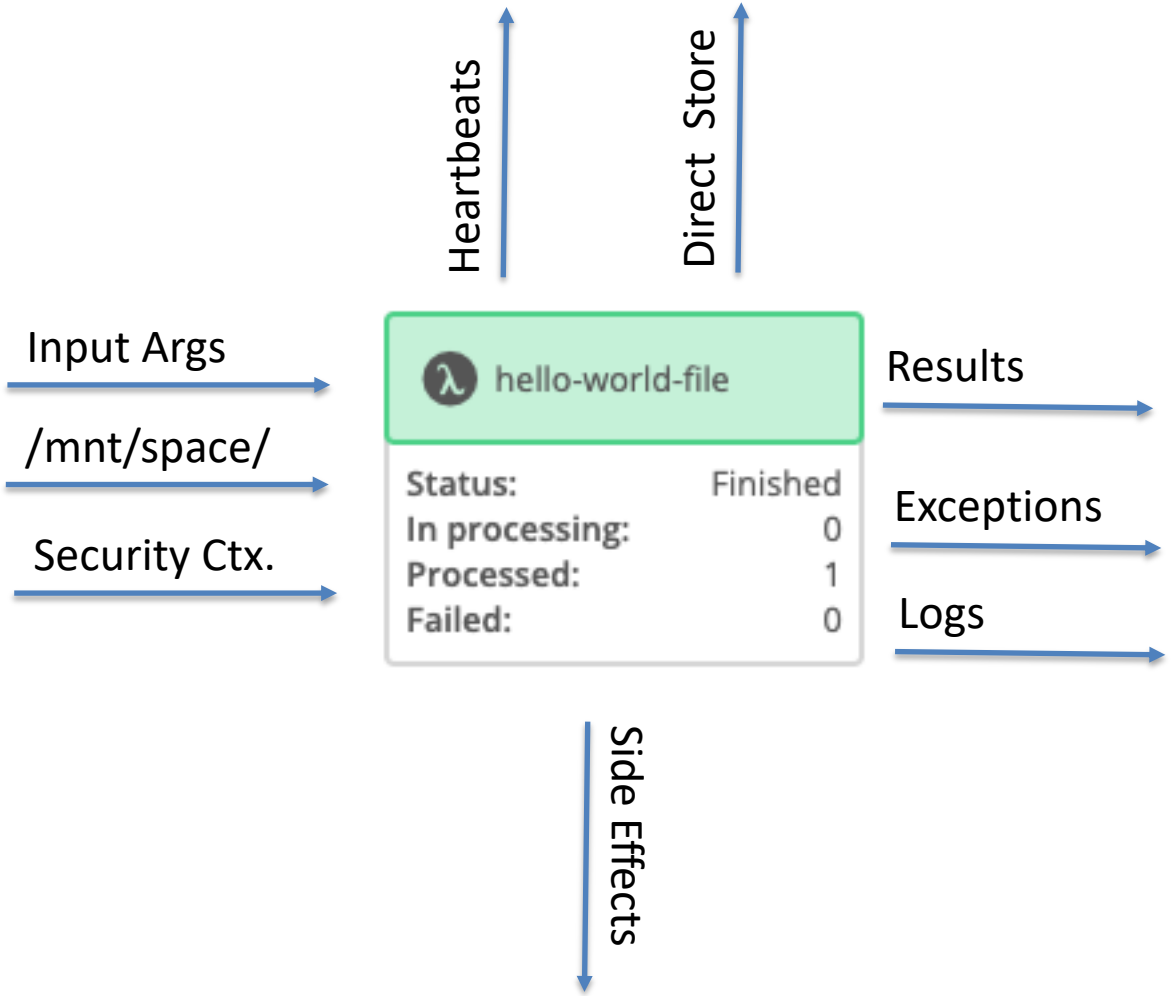

WORKFLOW

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

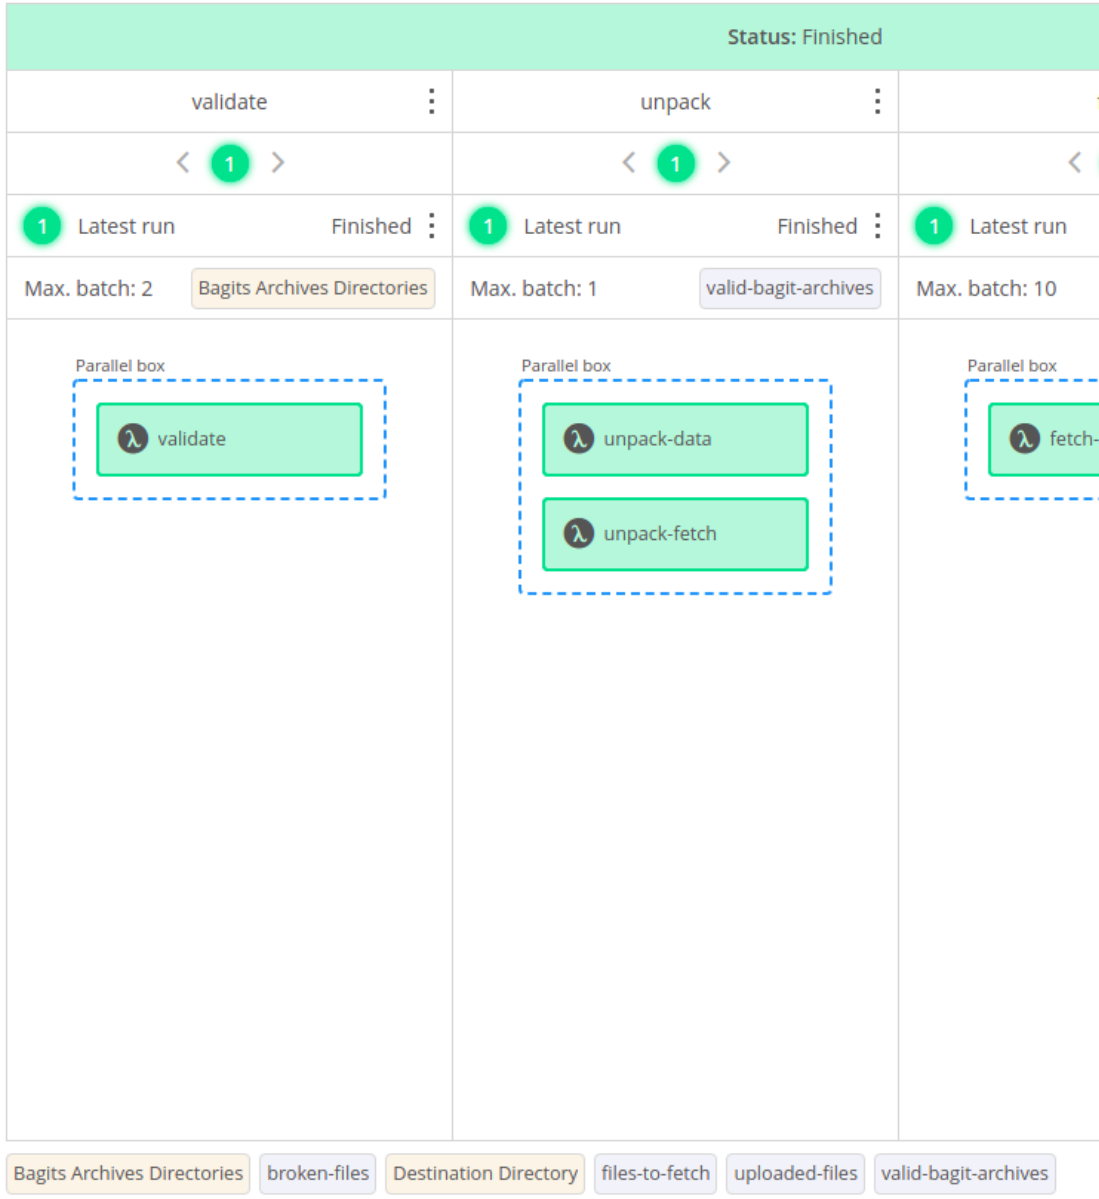

STORE

Store Anatomy

- **Persistent.** Keeps information to be iterated
- **Internal Model.** List, KV Map, Single Object, Forest Tree, Histogram for time series data
- **Strict Types.** One of: Object, File, AnyFile, Directory, String, etc.
- **Input User.** Defined before workflow execution.
- **Browsable**. User can see the current and saved status of all stores until the workflow execution is purged

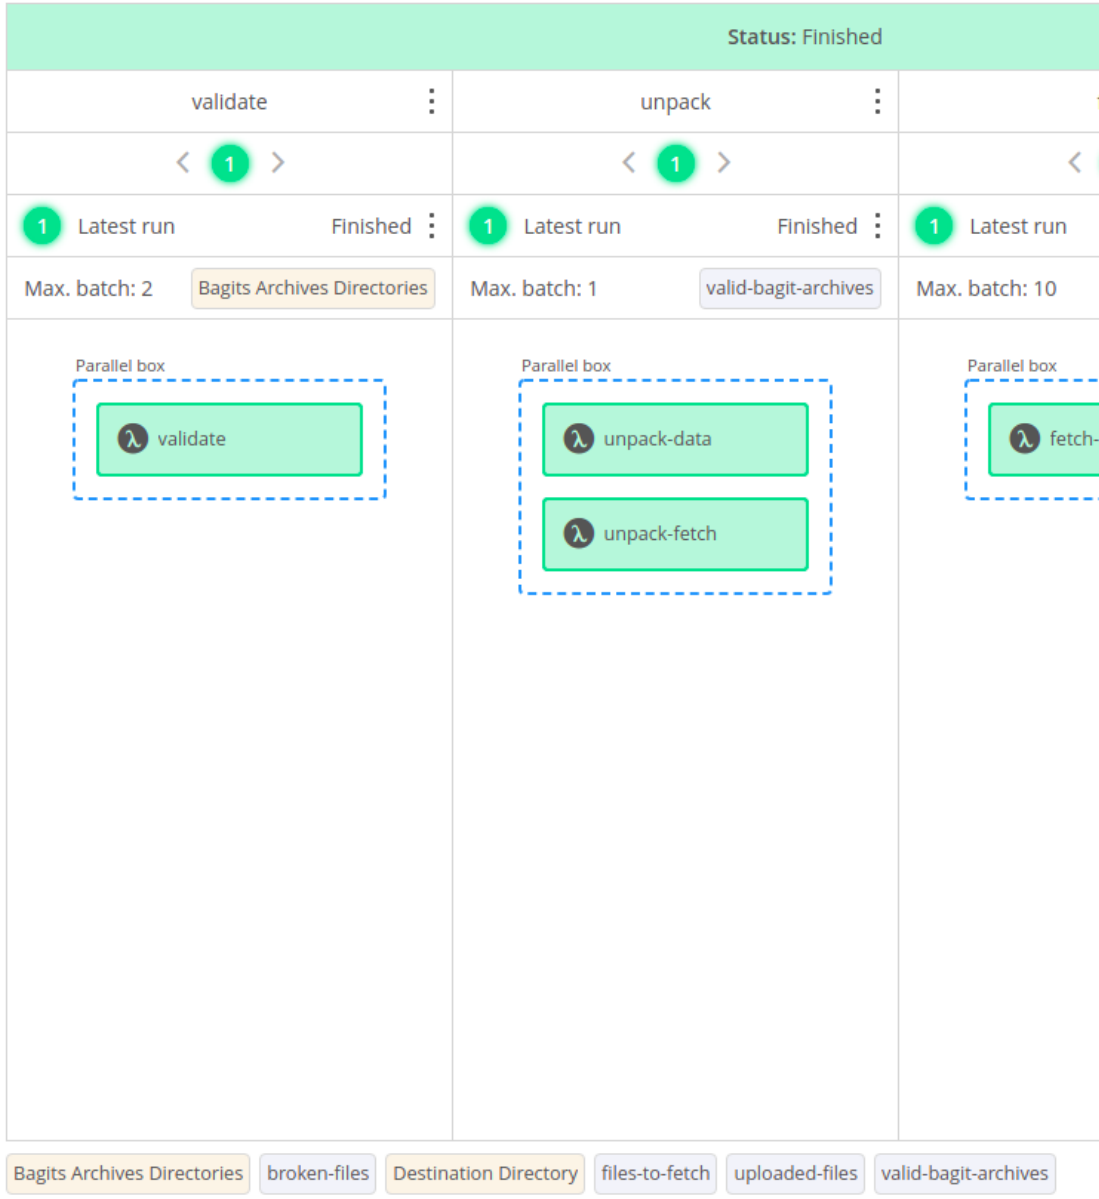

INVENTORY

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

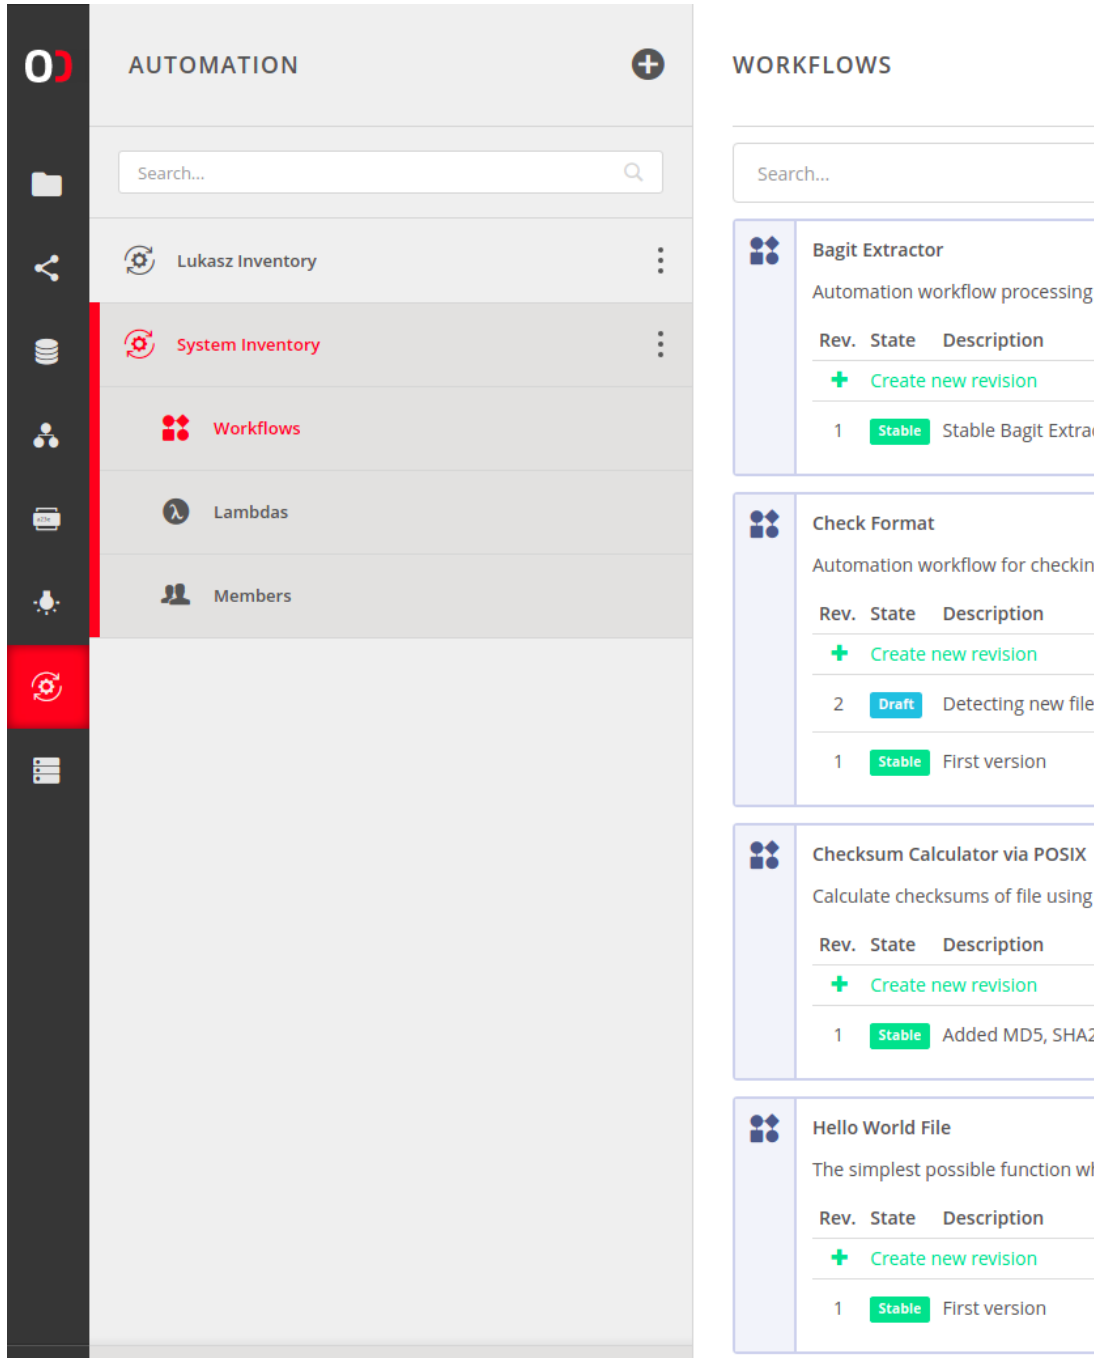

QUESTIONS?

Please visit:

www.onedata.org

EXAMPLE USECASES