Onedata FaaS Workflow Engine for Archivisation and beyond

Presented by: Lukasz Dutka

ONEDATA

In Collaboration With

T · · Systems ·

GWDG
Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

Open Archival Information System (OAIS)

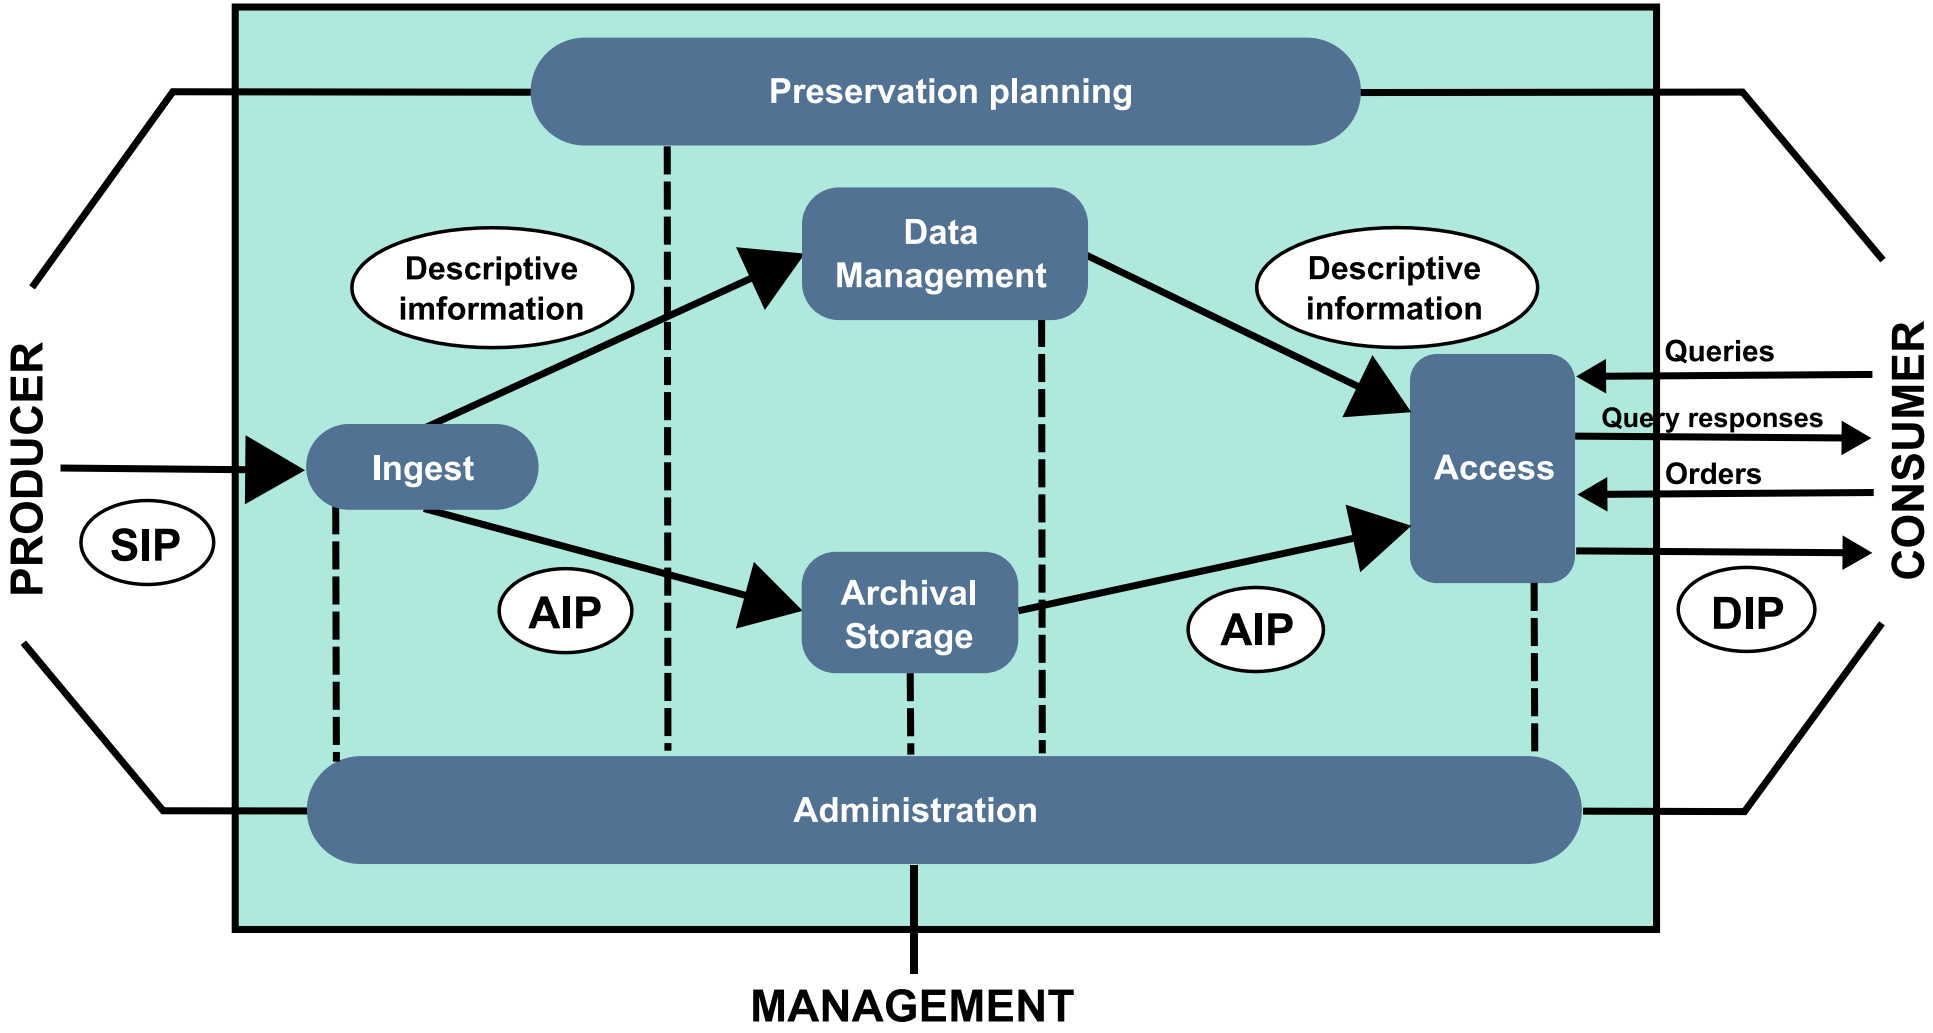

ARCHITECTURE

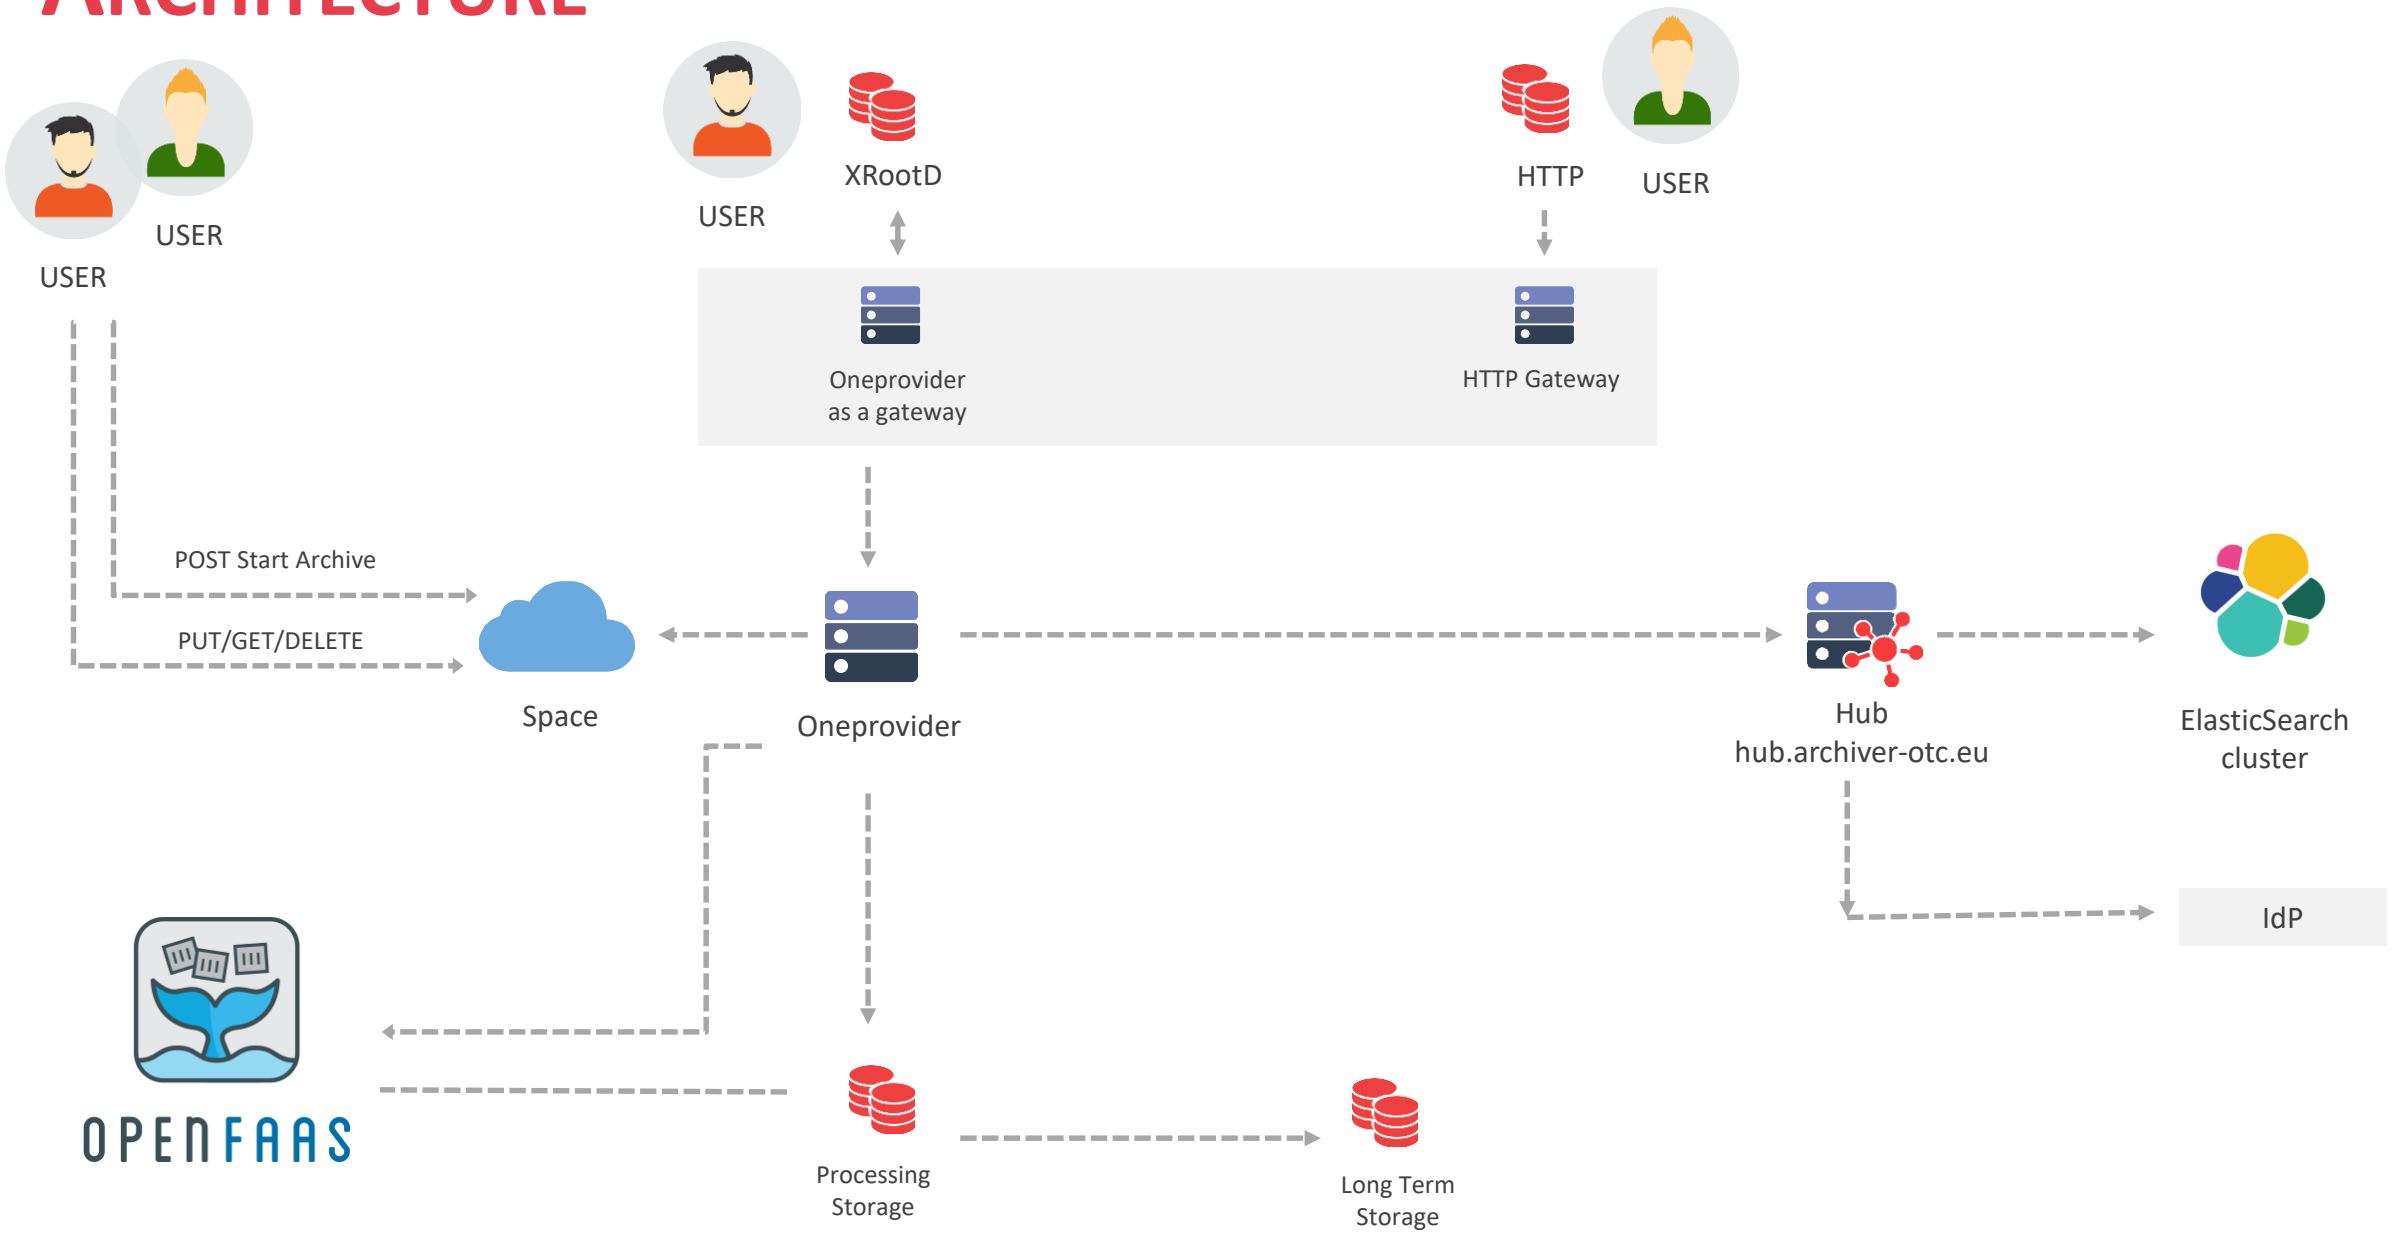

BAGIT FILES

Anatomy

- **SIP.** Popular format for submission packages.
- **Metadata.** Delivers metadata information about the content
- **Checksums.** Contains hashes of data and itself for consistency validation
- **Deliver Data.** Delivers data for smaller collections
- **List of Data to be fetched.** Describes the way to fetch (PULL) into the system
- **Standard?.** Unfortunately there many flavours of Bagits and several custom extensions.

```
$ cat fetch.txt
http://packages.devel.onedata.org/apt/ubuntu/1
02/pool/main/e/erlang/erlang-asn1_20.2.2+dfsg-
ubuntu1ppa6~xenial_amd64.deb 756358 data/dir1/
ir12/erlang-asn1_20.2.2+dfsg-0ubuntu1ppa6~xeni
l_amd64.deb
http://packages.devel.onedata.org/apt/ubuntu/1
02/pool/main/e/erlang/erlang-base-hipe_20.2.2+
fsg-0ubuntu1ppa6~xenial_amd64.deb 9136096 data
erlang-base-hipe_20.2.2+dfsg-0ubuntu1ppa6~xeni
l_amd64.deb
http://packages.devel.onedata.org/apt/ubuntu/1
02/pool/main/e/erlang/erlang-base_20.2.2+dfsg-
ubuntu1ppa6~xenial_amd64.deb 7386200 data/dir1
dir12/erlang-base_20.2.2+dfsg-0ubuntu1ppa6~xen
al_amd64.deb
http://packages.devel.onedata.org/apt/ubuntu/1
02/pool/main/e/erlang/erlang-common-test_20.2.
+dfsg-0ubuntu1ppa6~xenial_amd64.deb 1065320 da
a/dir1/erlang-common-test_20.2.2+dfsg-0ubuntu1
pa6~xenial_amd64.deb
```

New Features

Automation Engine

01

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**

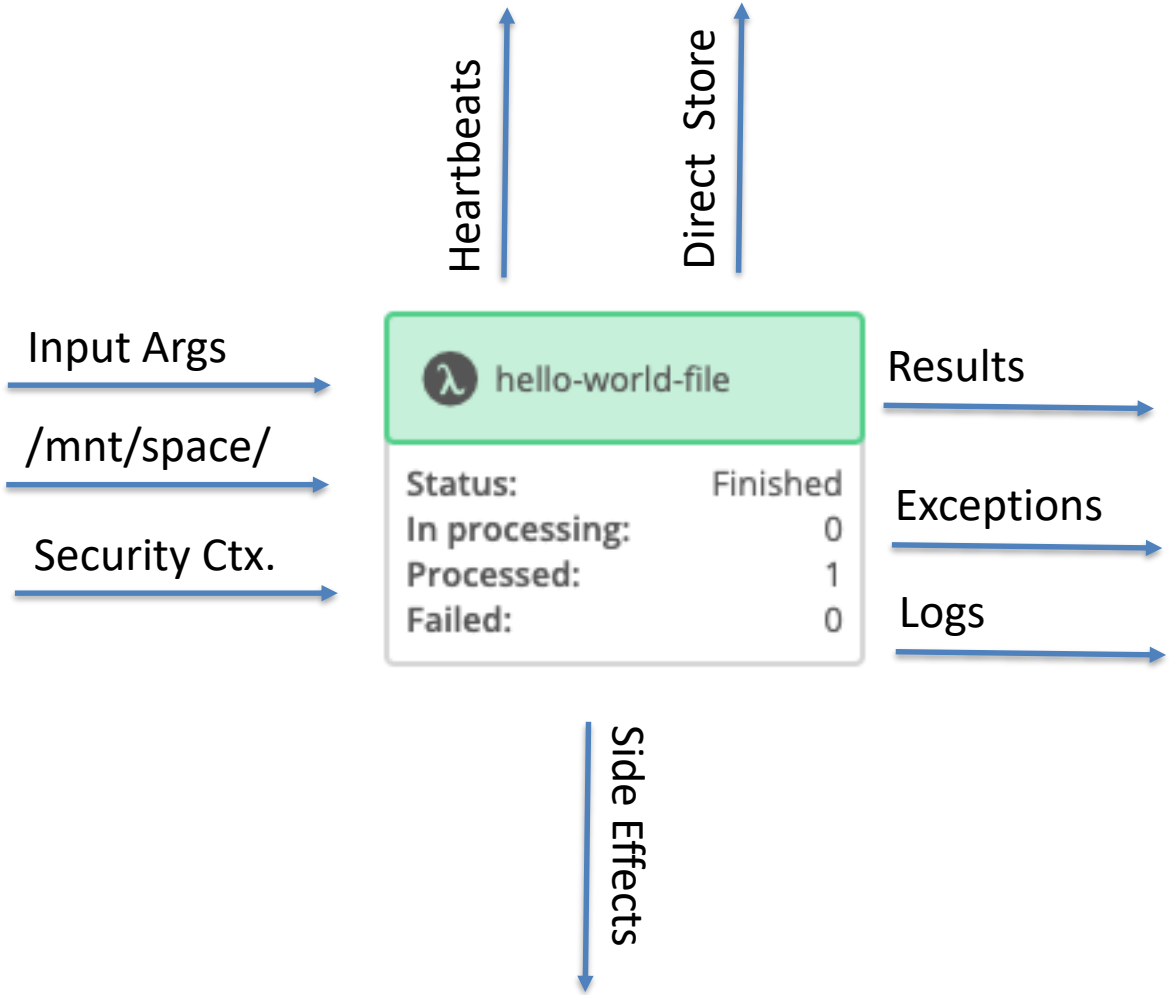

WORKFLOW

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

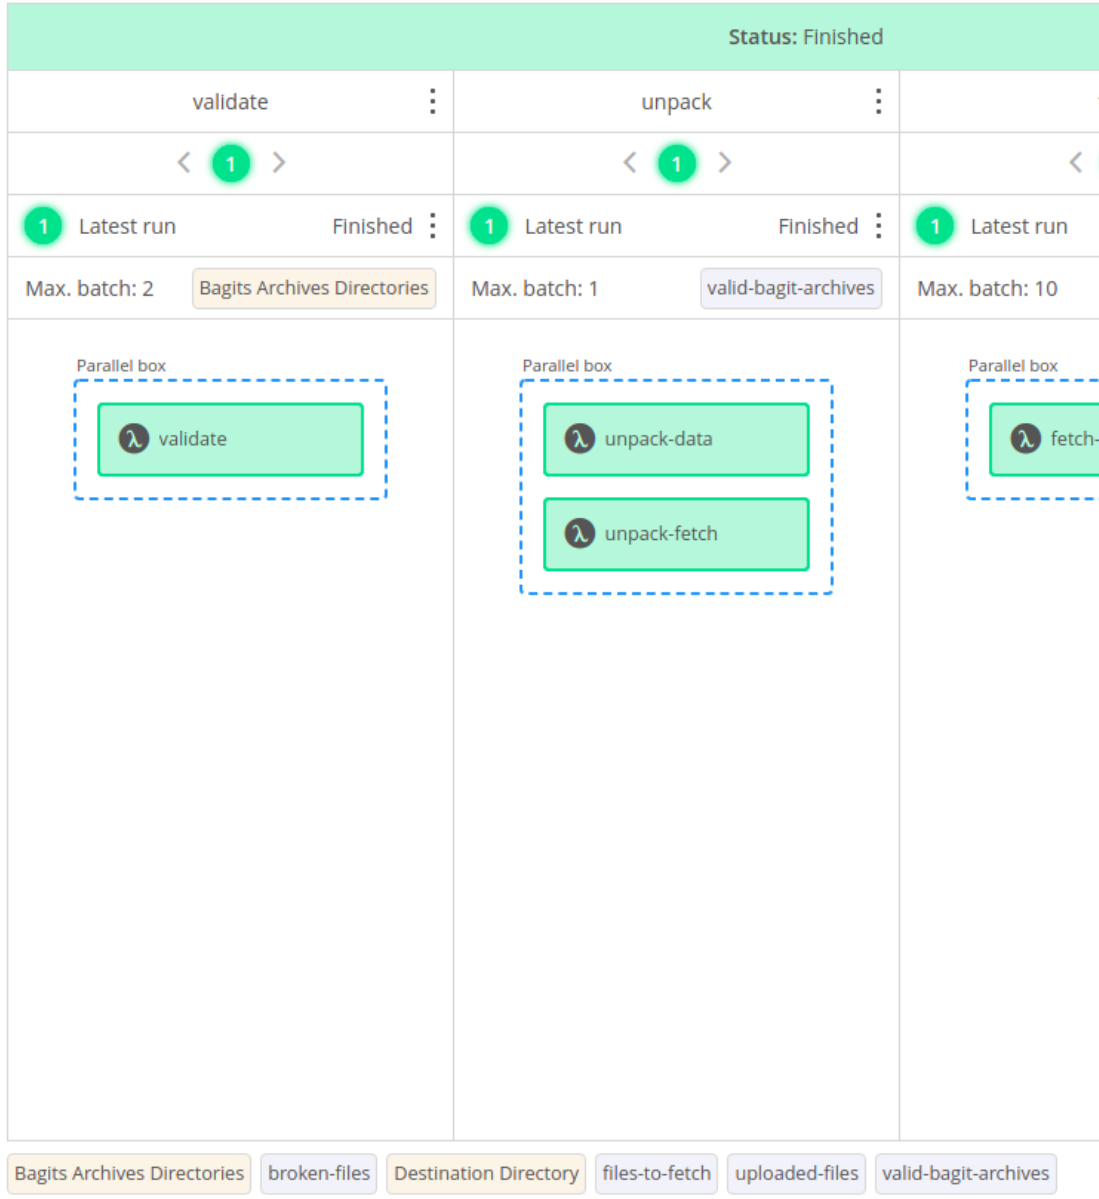

STORE

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

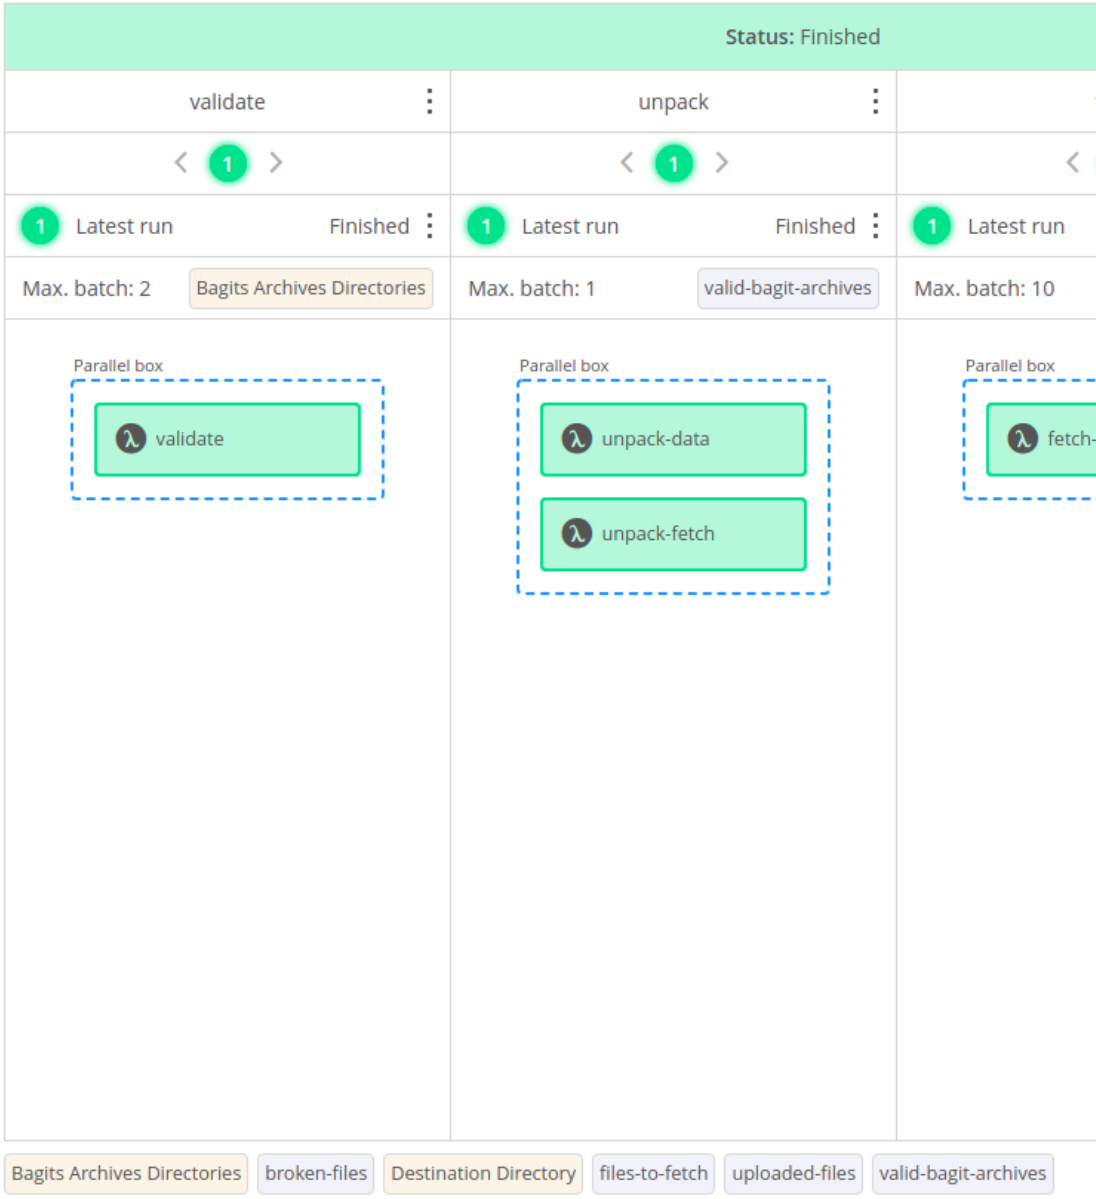

INVENTORY

Inventory Anatomy

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

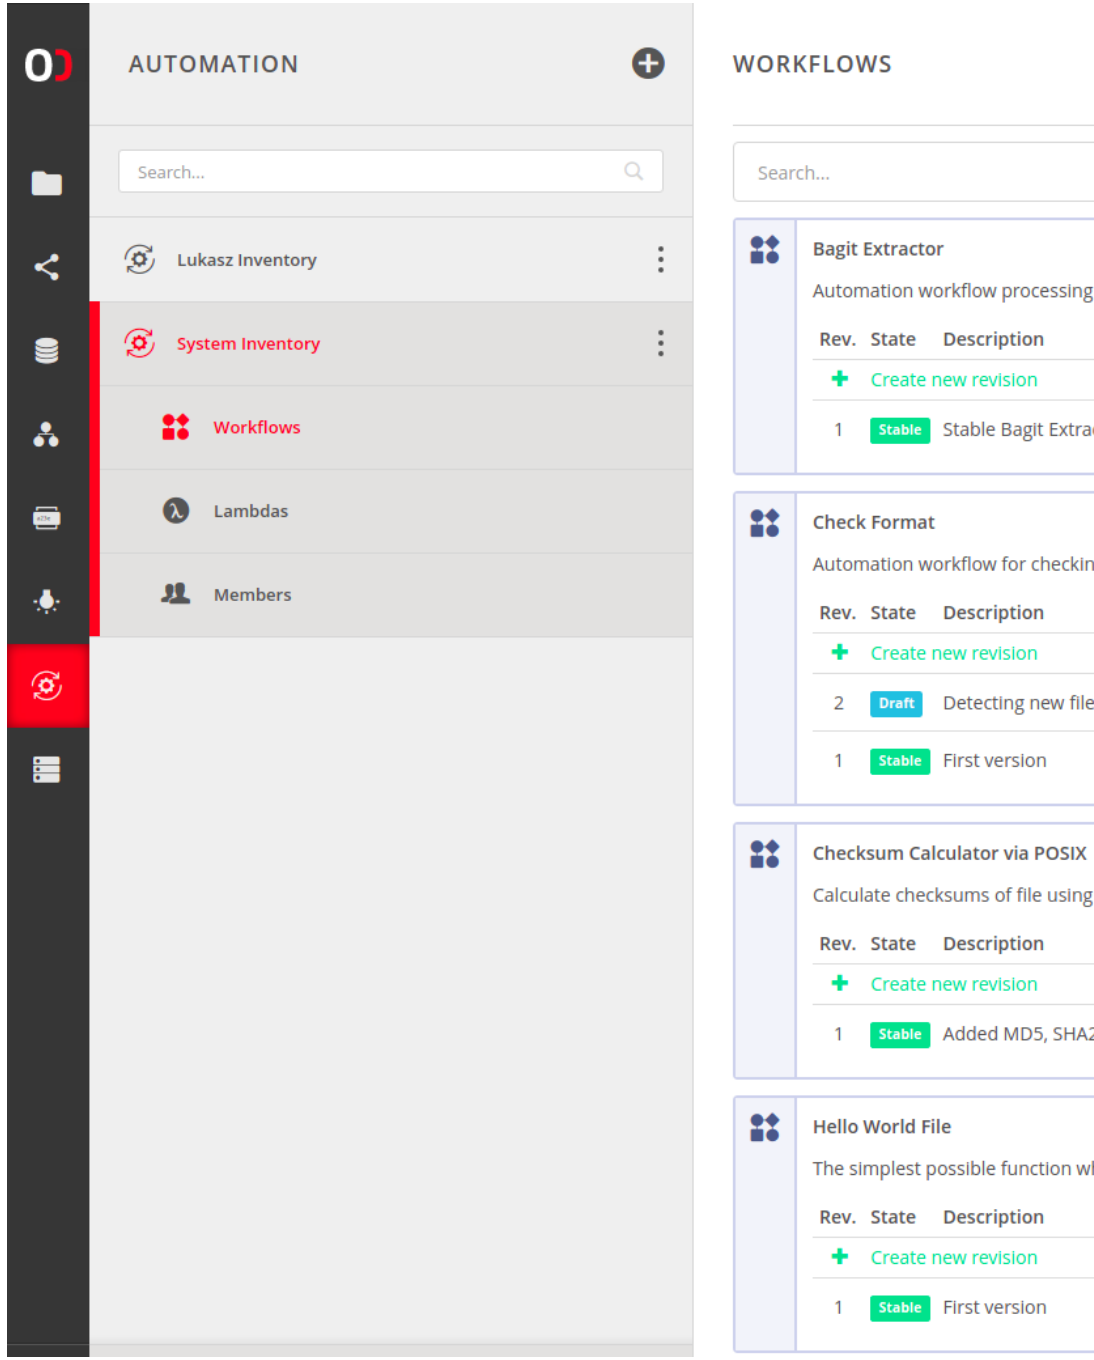

DEMO

New Features

OAIS Archivisation

02

Datasets, AIP, DIP, Incremental Archives

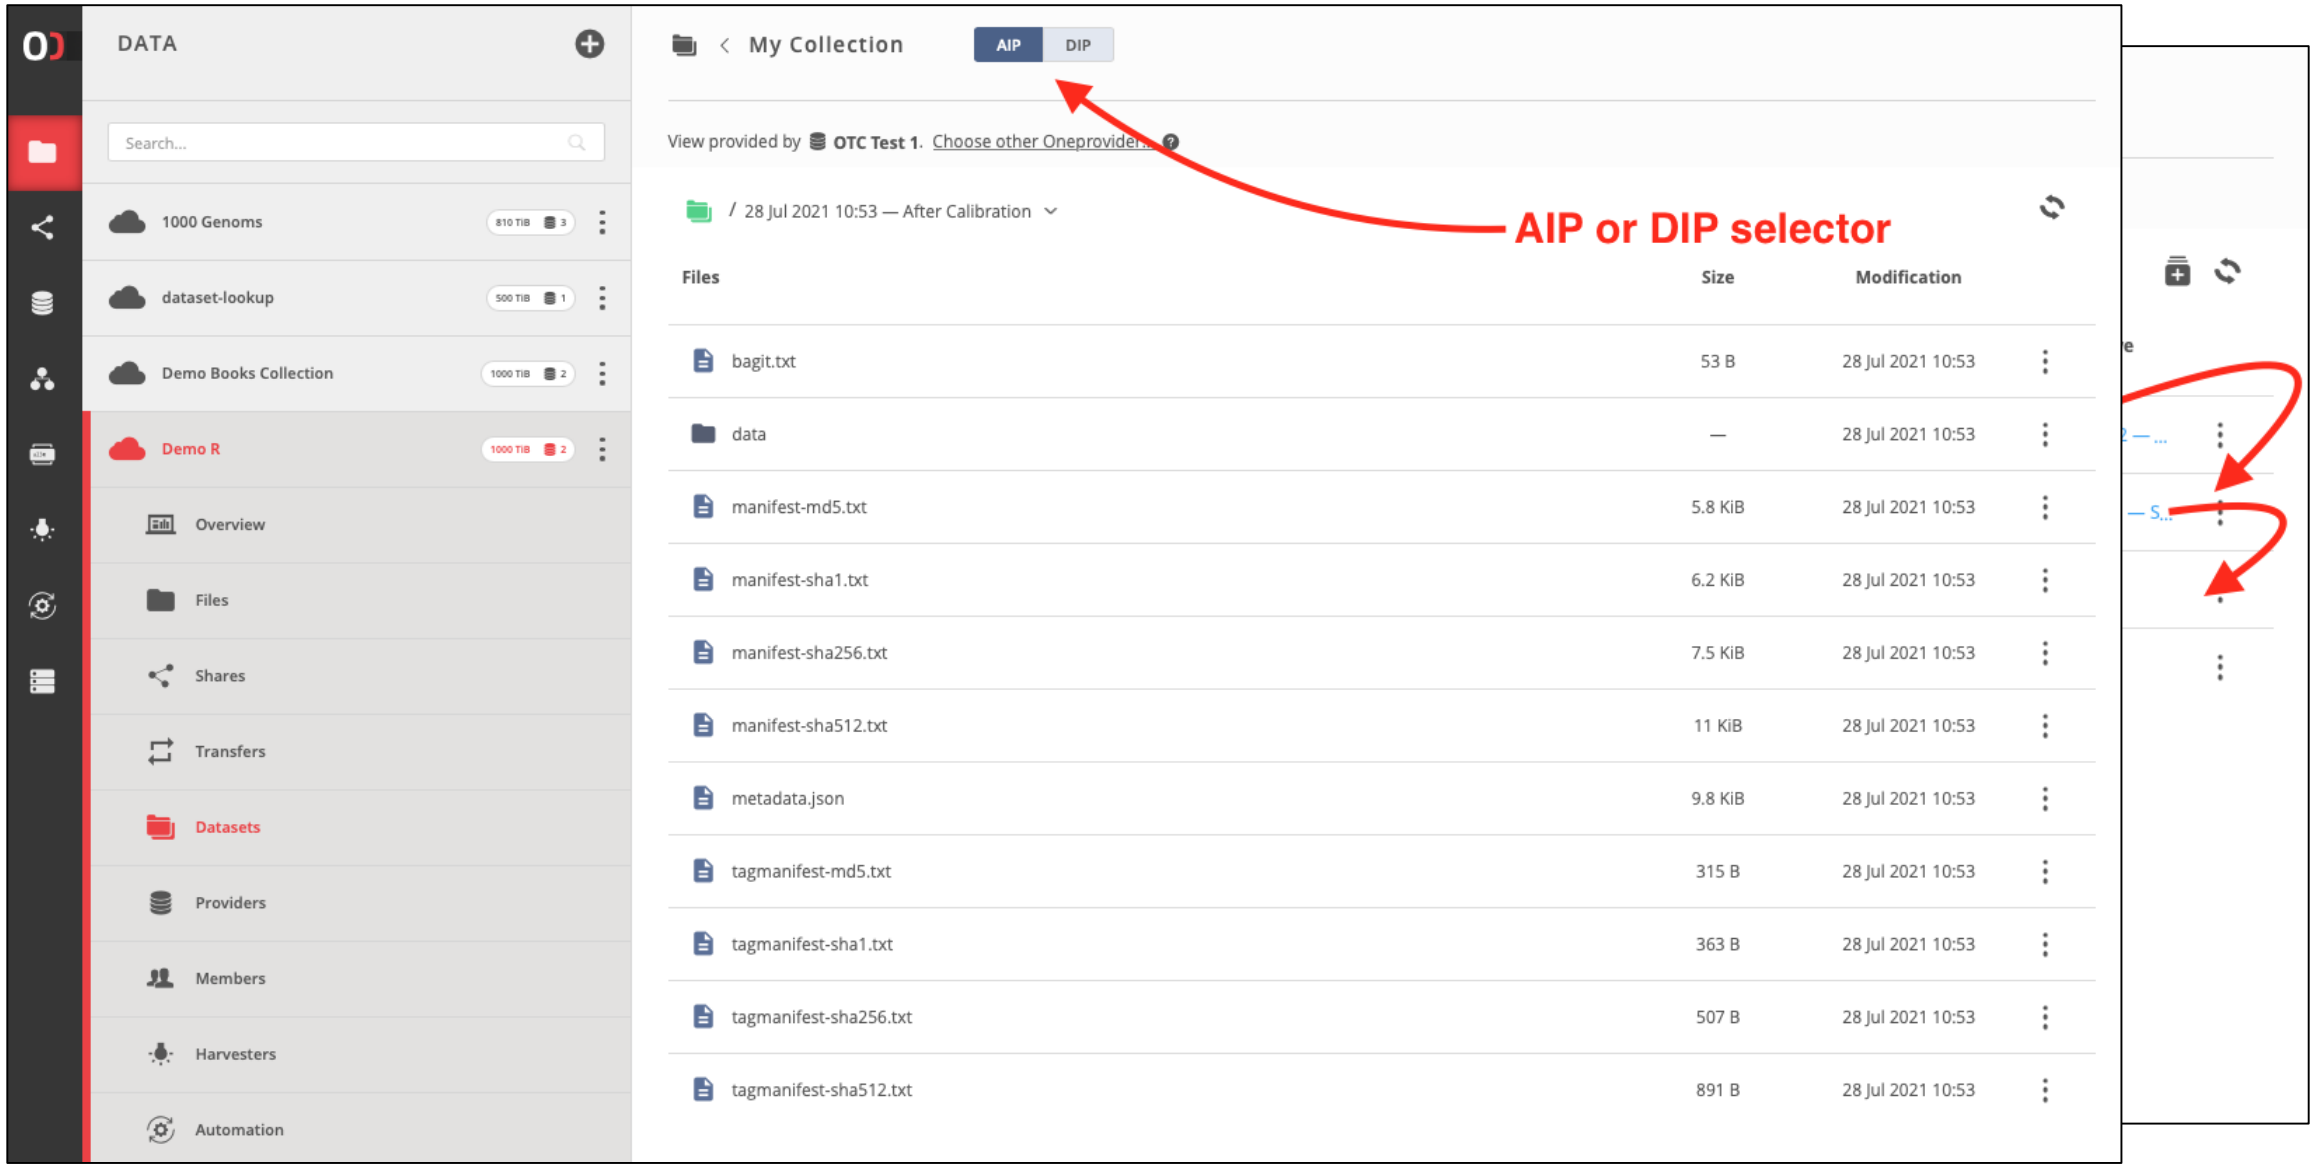

New Features

Hierarchical Datasets and Archives

03

Datasets Hierarchy – Embedded Archives

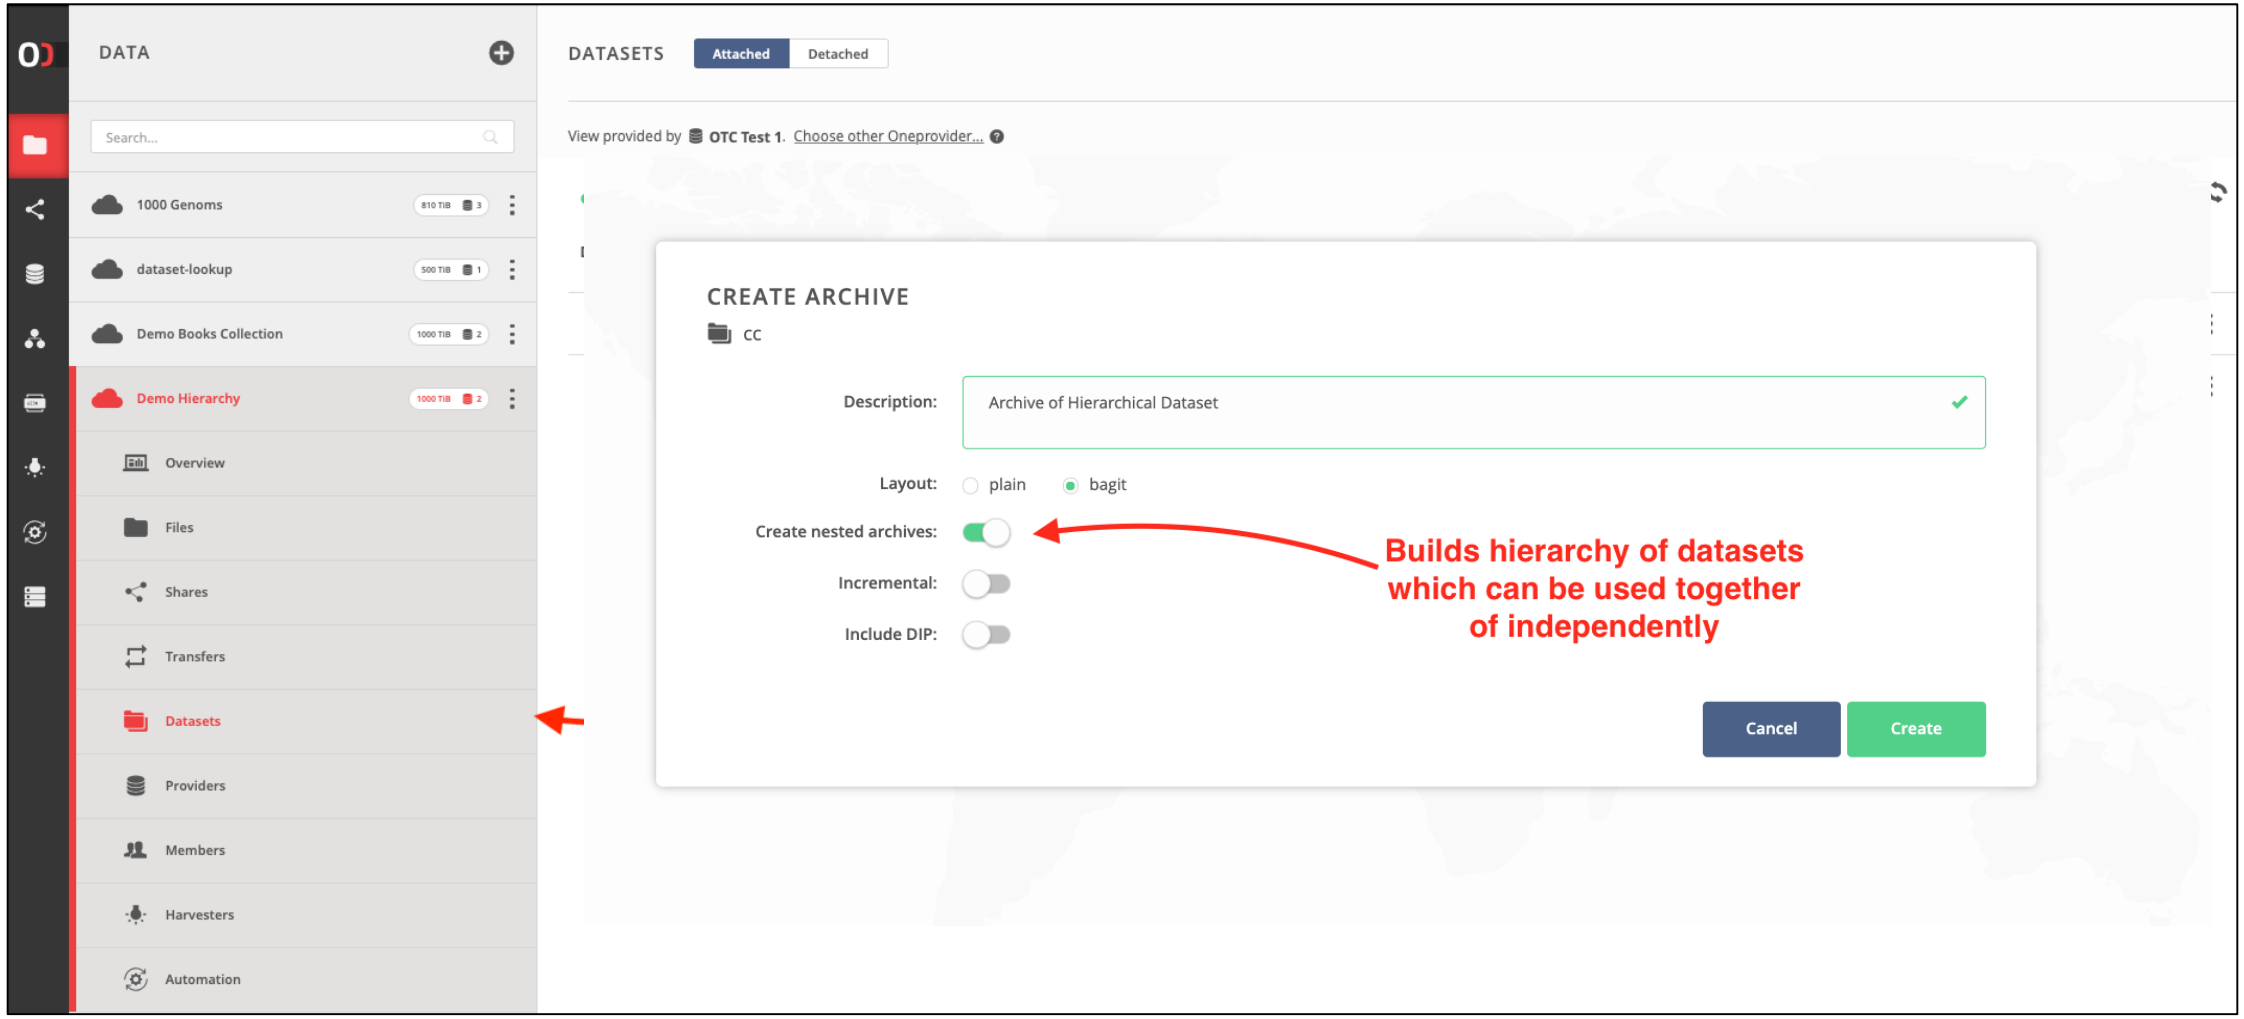

New Features

Hard links and Symbolic links

04

Hard Links

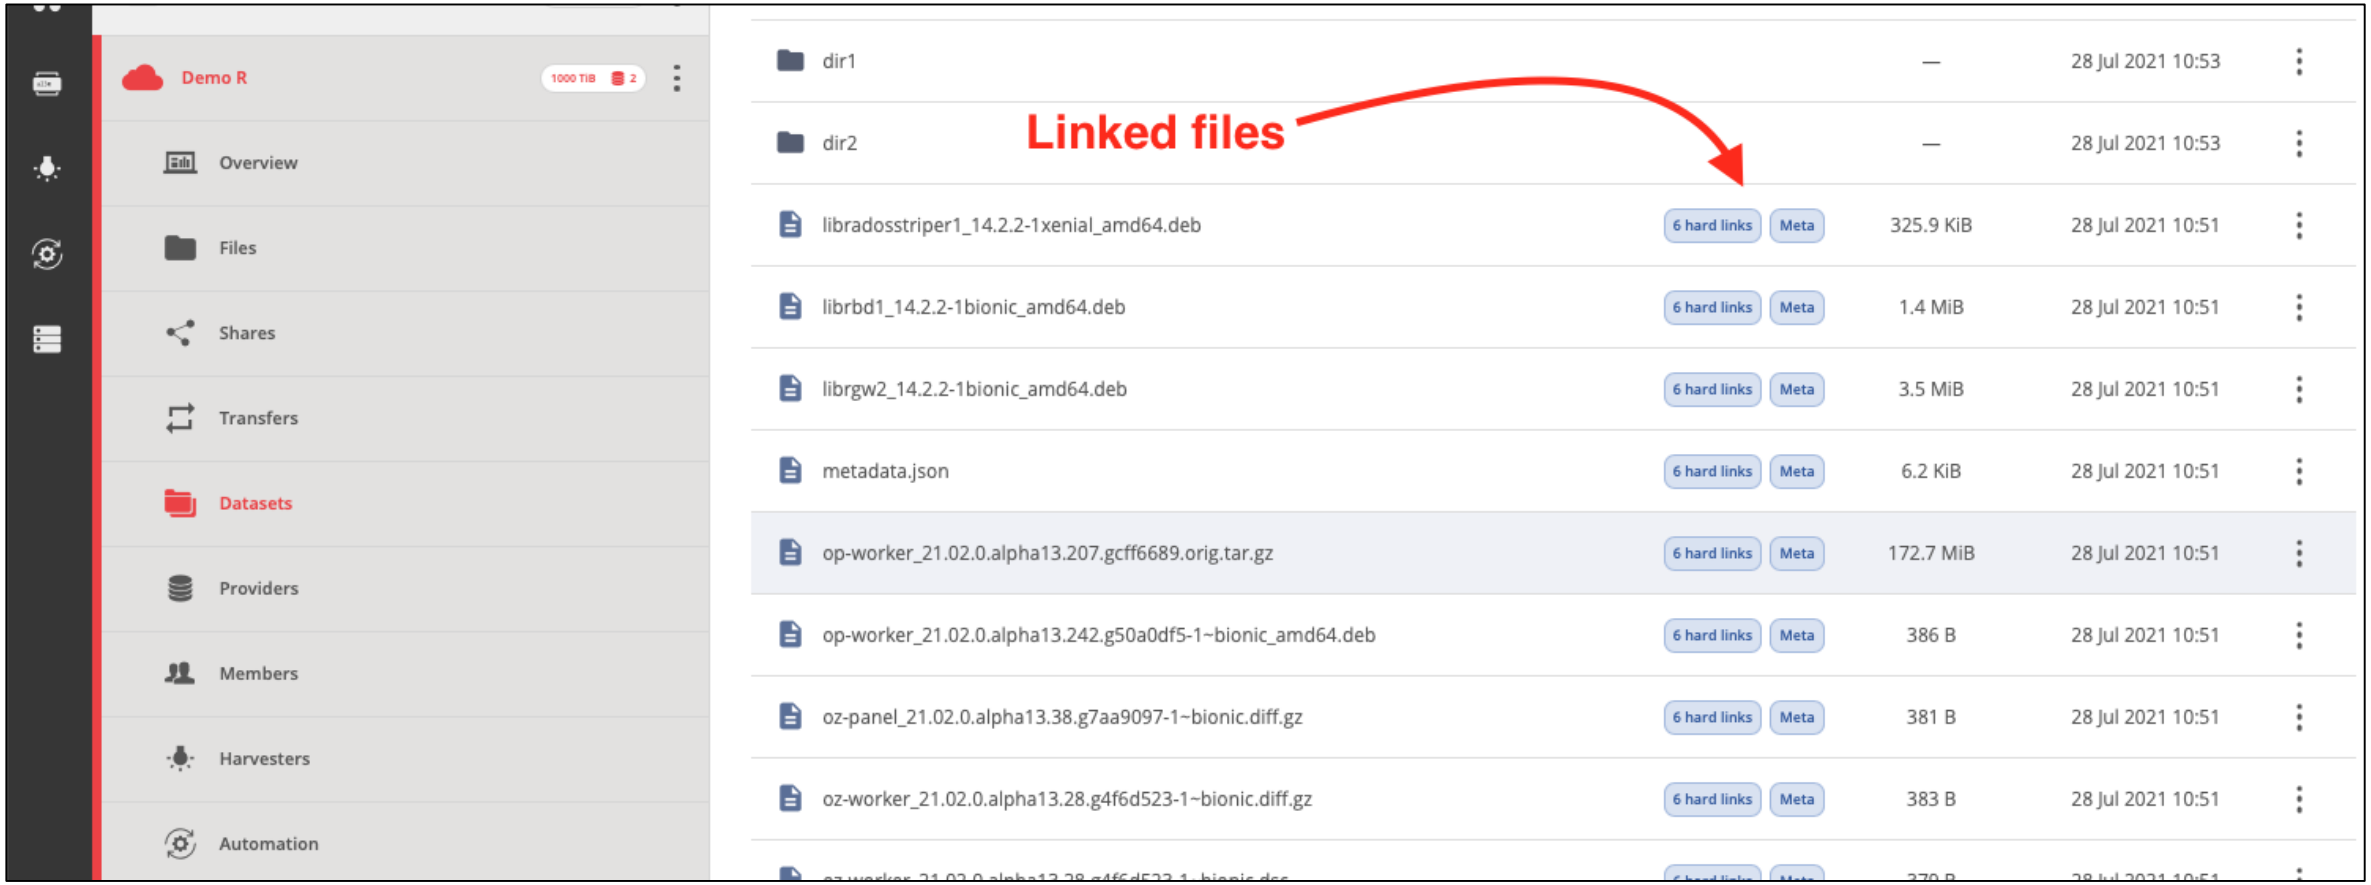
Demo R
1000 TiB
2
Overview
Files
Shares
Transfers
Datasets
Providers
Members
Harvesters
Automation
Linked files
dir1 — 28 Jul 2021 10:53
dir2 — 28 Jul 2021 10:53
libradosstriper1_14.2.2-1xenial_amd64.deb 6 hard links Meta 325.9 KiB 28 Jul 2021 10:51
librbd1_14.2.2-1bionic_amd64.deb 6 hard links Meta 1.4 MiB 28 Jul 2021 10:51
librgw2_14.2.2-1bionic_amd64.deb 6 hard links Meta 3.5 MiB 28 Jul 2021 10:51
metadata.json 6 hard links Meta 6.2 KiB 28 Jul 2021 10:51
op-worker_21.02.0.alpha13.207.gcff6689.orig.tar.gz 6 hard links Meta 172.7 MiB 28 Jul 2021 10:51
op-worker_21.02.0.alpha13.242.g50a0df5-1~bionic_amd64.deb 6 hard links Meta 386 B 28 Jul 2021 10:51
oz-panel_21.02.0.alpha13.38.g7aa9097-1~bionic.diff.gz 6 hard links Meta 381 B 28 Jul 2021 10:51
oz-worker_21.02.0.alpha13.28.g4f6d523-1~bionic.diff.gz 6 hard links Meta 383 B 28 Jul 2021 10:51

Symbolic Links

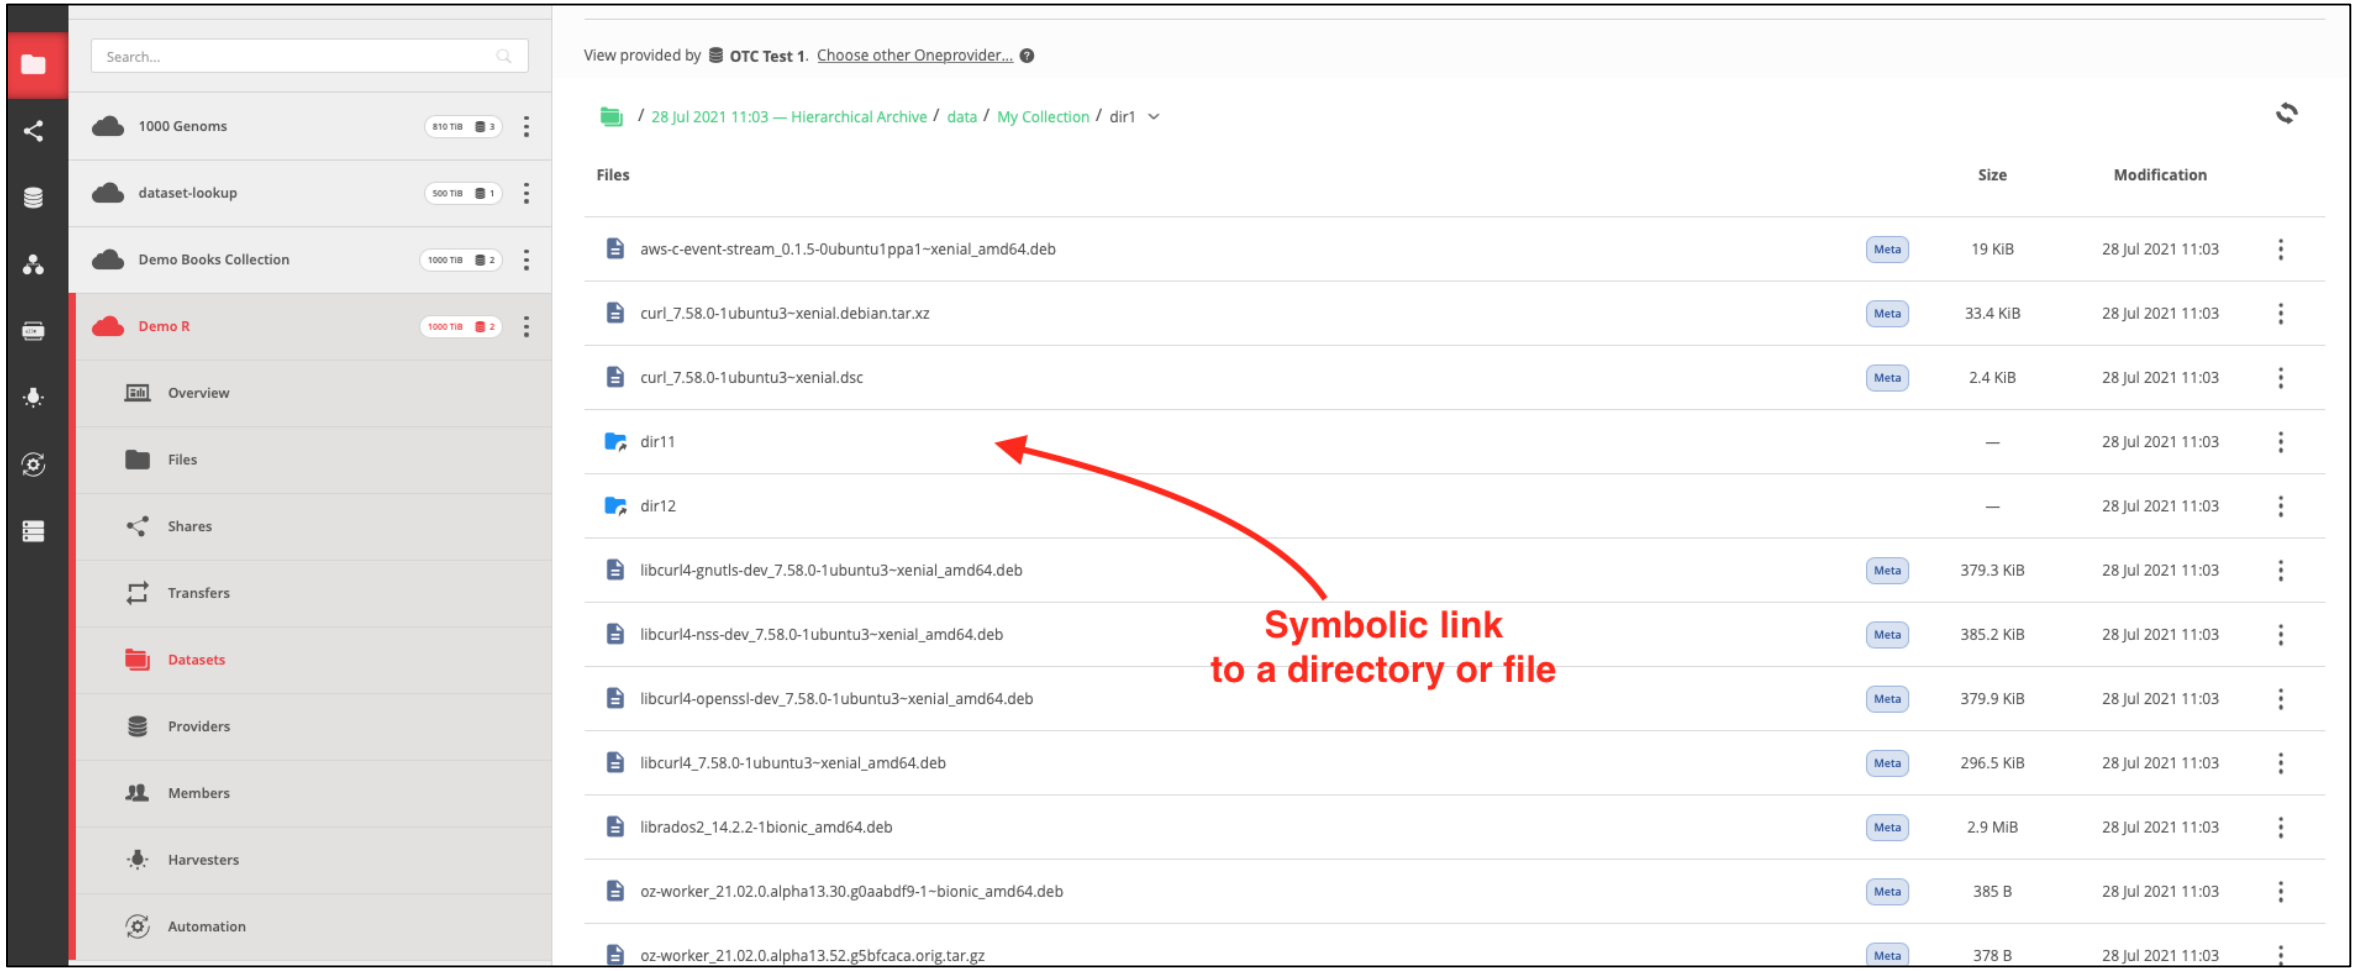

R&D Achievements

Performance Improvements

05

Push Ingestion Speed

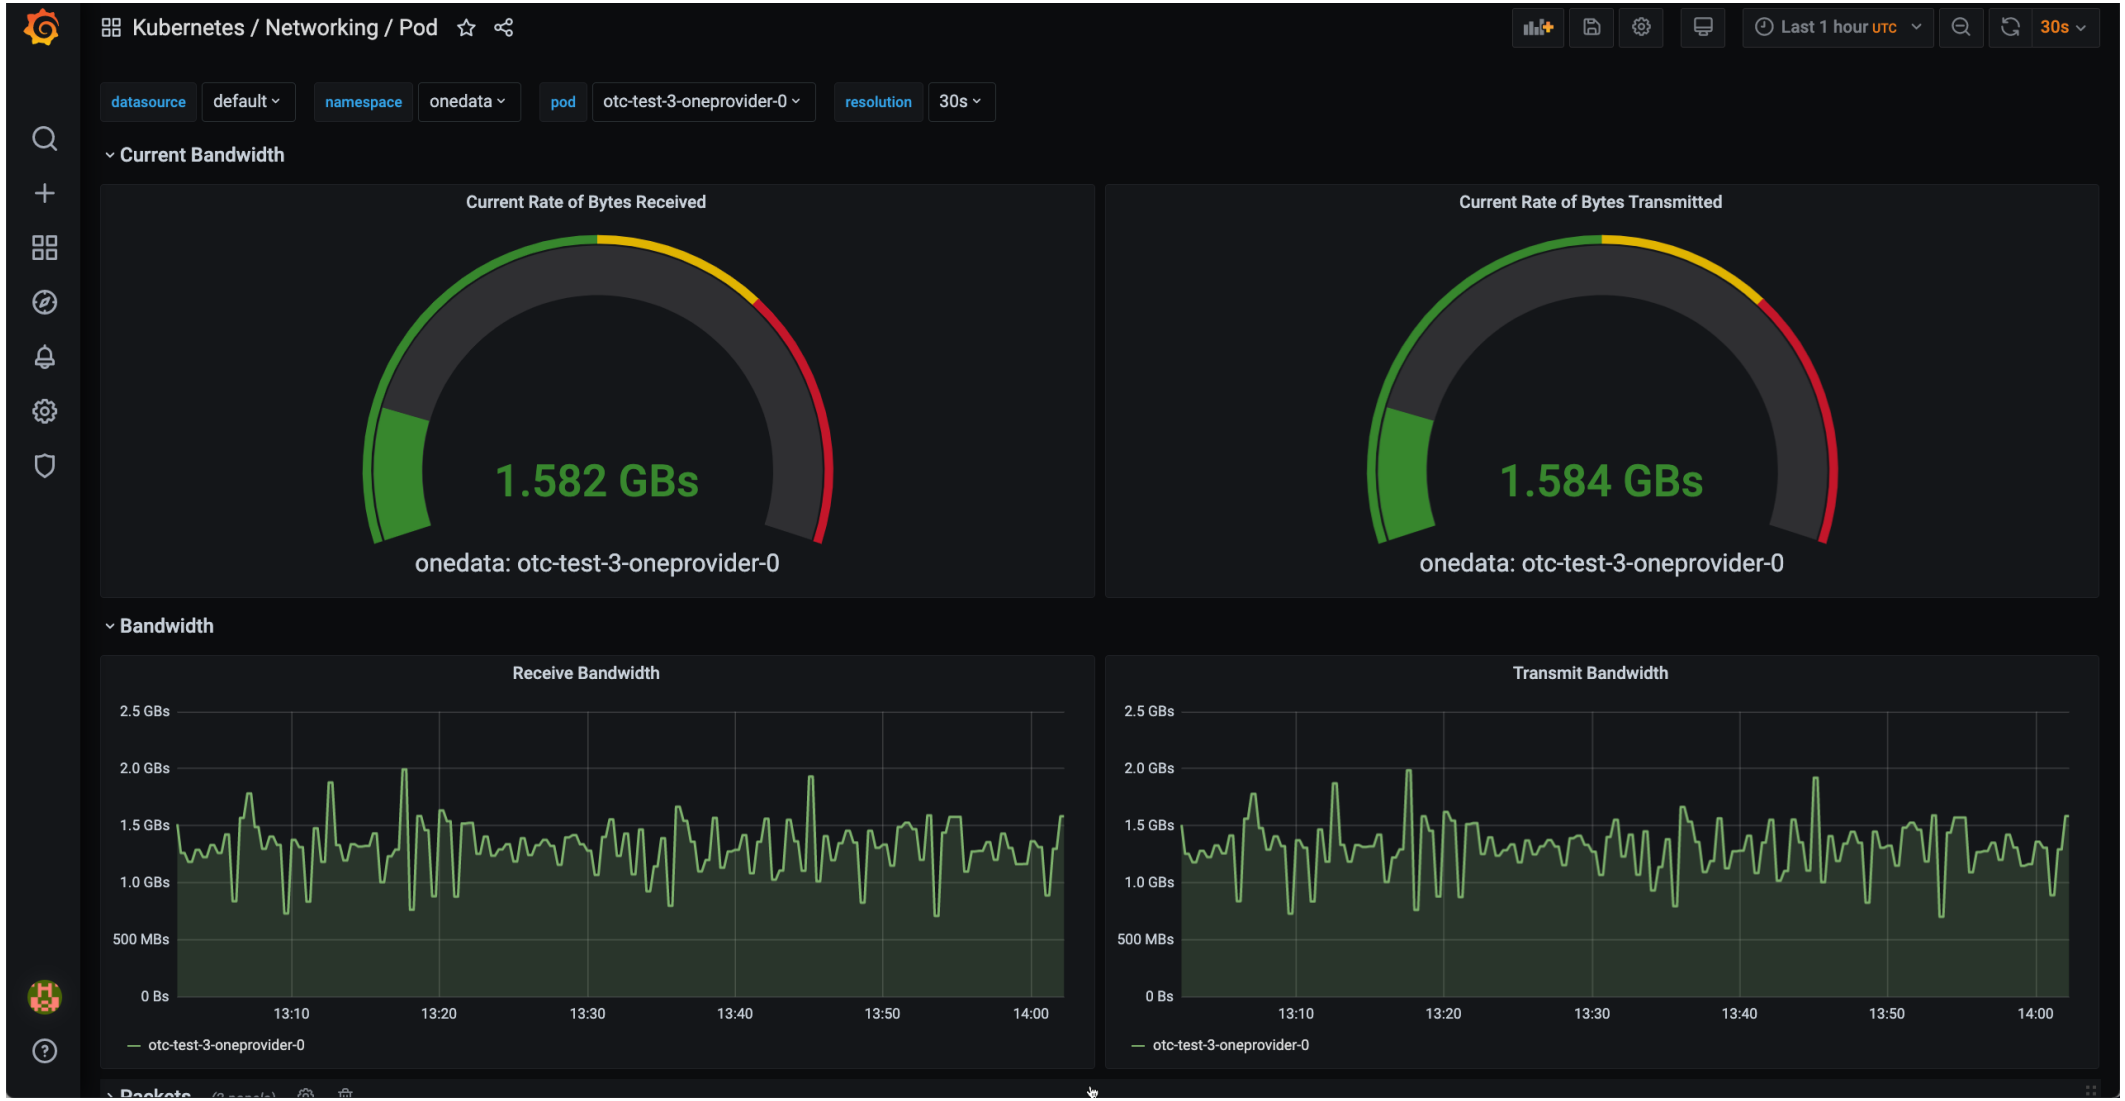

Benefits of Our Approach

Data Mgmt. & Discovery

- Built around **hybrid** and distributed architecture
- Support well known concepts: **files, folders, links, symbolic links, datasets, archives, metadata**
- Harvesting and complex **indexing** for data **discovery**
- Deep integration with IdPs and multi-level **access control**
- **PUSH** and **PULL** ingestion
- **Open Data** ready

Data Processing

- Virtual File System **POSIX** API
- Automation Workflows
- Flexible **automation functions** driven by community and end-users
- Out of the box **auto-scalability**
- High **throughput** performance
- No need for expensive temporary POSIX block storage

Licensing & Platform

- Fully Open Source - MIT
- Platform Agnostic
- OpenFaaS support
- Kubernetes ready
- Backend Storage Agnostic

Selected Future Plans

- Storage grants distributing capacity from procurer to end-user
- UX Improvements e.g. External storage at the user level
- Enhancing performance to the 100 Gbps level
- CTS certification
- Multi-tiered storage management including tape storage
- OAIS Archival Wizard enabling lower entry barrier
- Data Lifecycle Polices

ONEDATA

QUESTIONS?