
# First thoughts about a Data Lake prototype

**Aris Fkiaras**

CERN, Geneva, Switzerland

Amsterdam, 4 July 2019

ESCAPE - The European Science Cluster of Astronomy & Particle Physics ESFRI Research Infrastructures has received funding from the European Union's Horizon 2020 research and innovation programme under the Grant Agreement n° 824064.

# Disclaimer

This is a **proposal**

Reference implementation, nothing is set in stone

Hopefully starting point for the discussion


Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# Why Datalake?

- Use an ecosystem of distributed and diverse storage solutions
- Have a distributed storage system for latency hiding and data integrity
- Manage file lifecycle, number and place of replicas

Requirements

- Hide internal mechanics from user


Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# ESCAPE Data infrastructure for Open Science

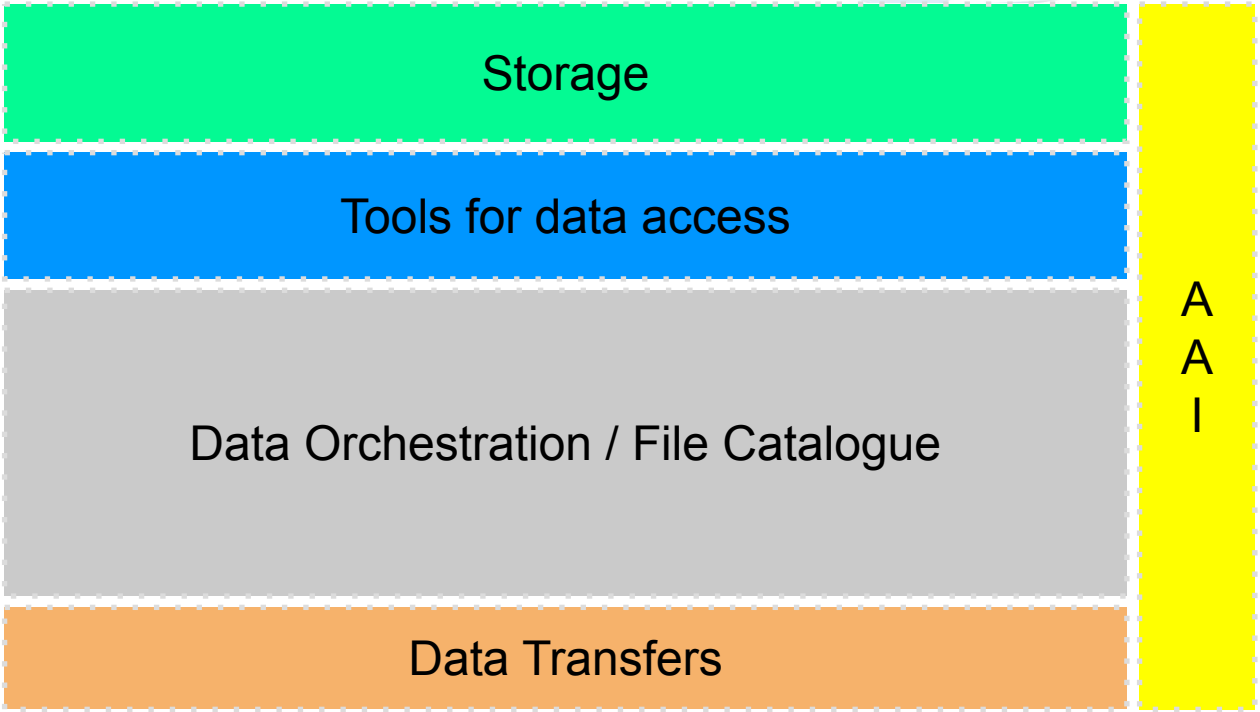



Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# ESCAPE Data infrastructure for Open Science

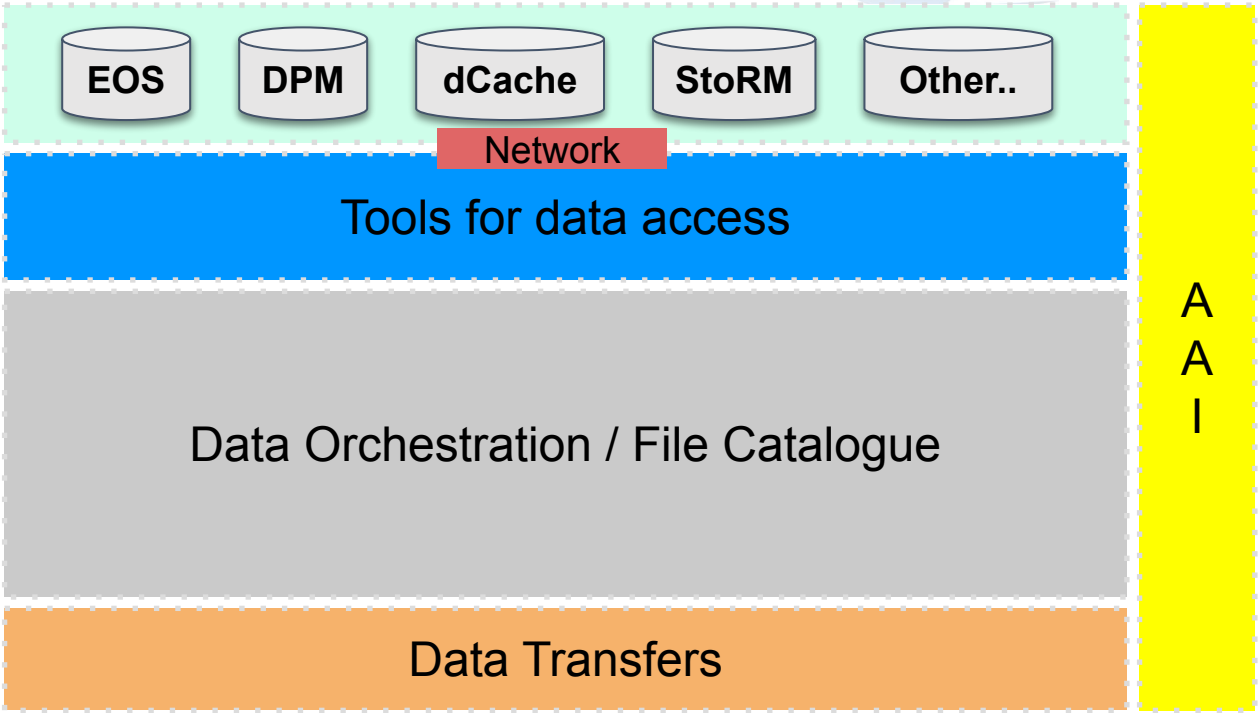



Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# ESCAPE Data infrastructure for Open Science

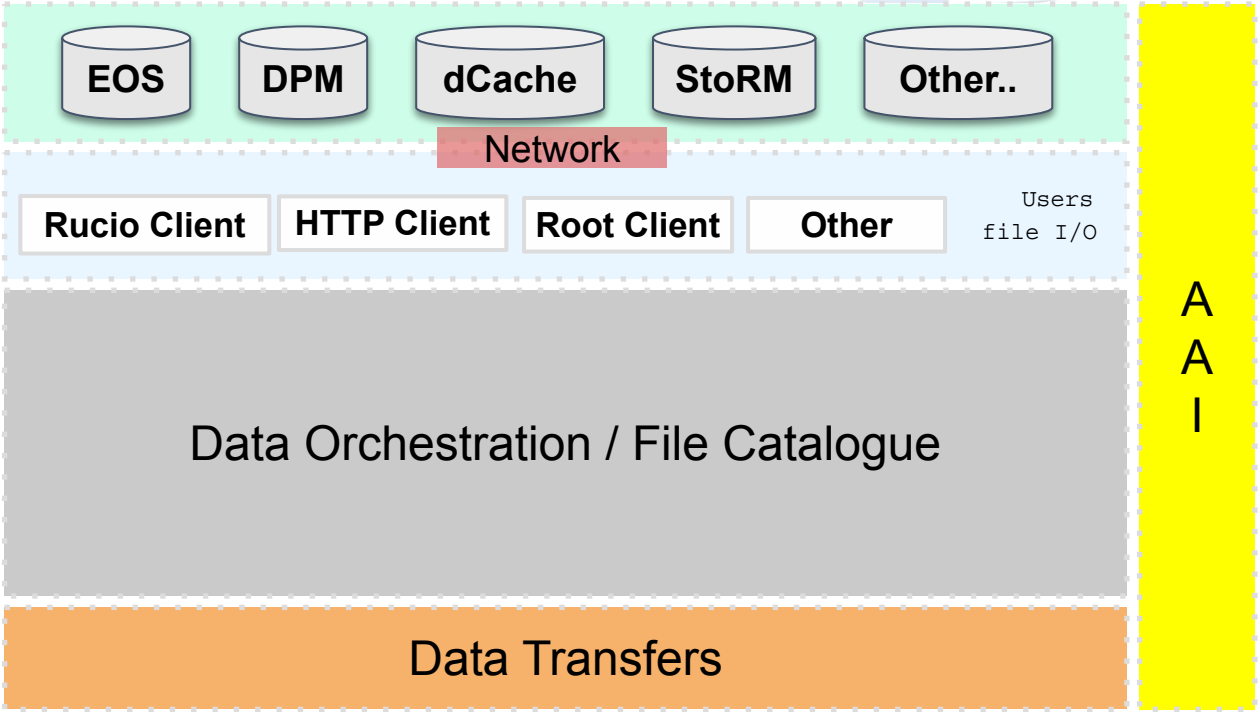



Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# ESCAPE Data infrastructure for Open Science

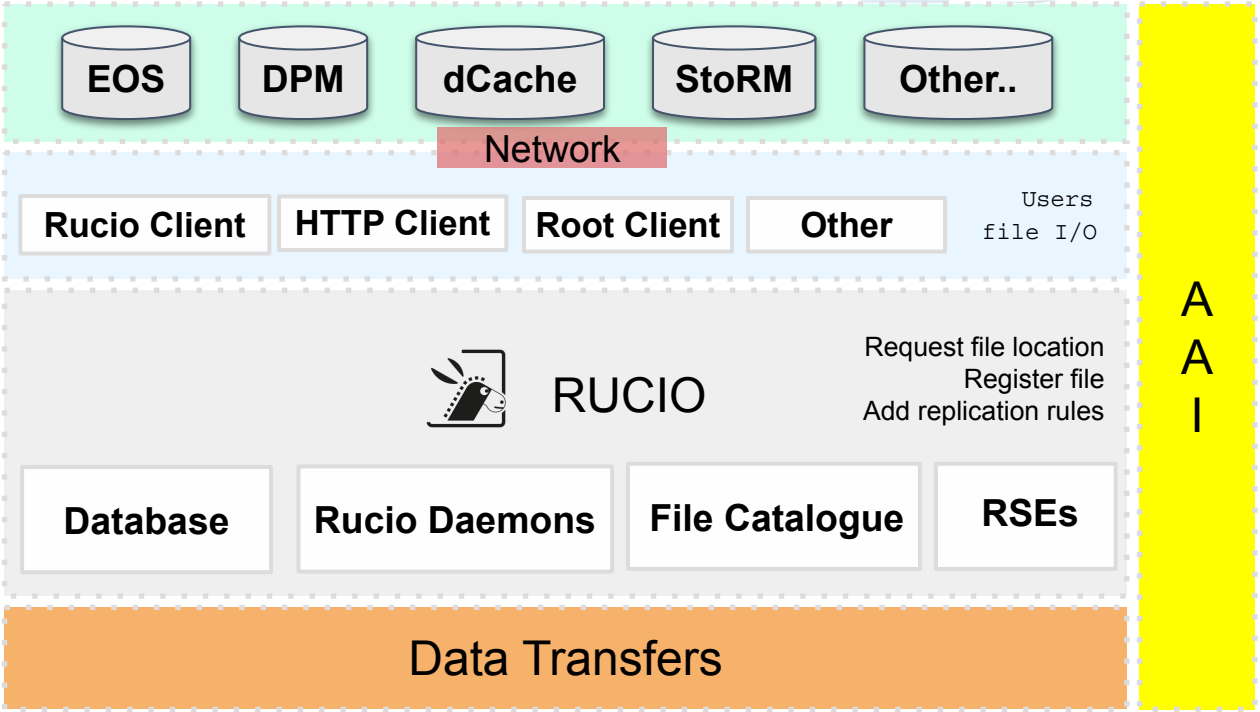



Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# RSEs

- A Rucio Storage Element (RSE) is the logical abstraction of a storage system for physical files.
- To set up a new RSE we need:
  - Hostname:Port
  - Scheme (root, http, gsiftp, ….)
  - Prefix

**EU-R-1.json** 880 Bytes

```
{
    "EU-R-1":{
        "backend_type": "XROOTD",
        "deterministic": true,
        "protocols": {
        "default": "root",
        "supported": {
            "root": {
            "prefix": "/eos/eulake/tests/rucio_test/eu_r_1",
            "impl": "rucio.rse.protocols.gfal.Default",
            "hostname": "eulake.cern.ch",
            "scheme": "root",
            "port": 1094,
            "domains": {
                "lan": {
                "read": 1,
                "write": 1,
                "delete": 1,
                "third_party_copy": 1
                },
                "wan": {
                "read": 1,
                "write": 1,
                "delete": 1,
                "third_party_copy": 1
                }
            }
            }
        }
        },
        "storage_usage_tool": "rucio.rse.protocols.xrootd.Default.getSpace"
    }
}
```


Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# ESCAPE Data infrastructure for Open Science

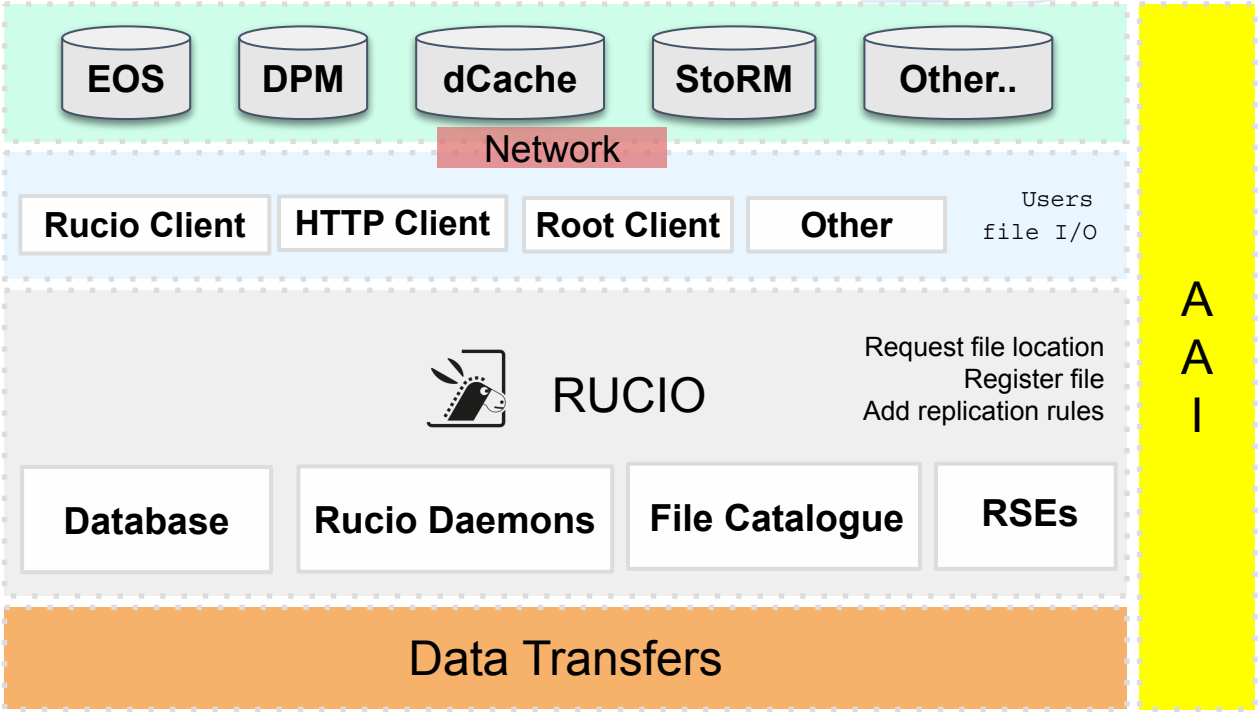



Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# ESCAPE Data infrastructure for Open Science

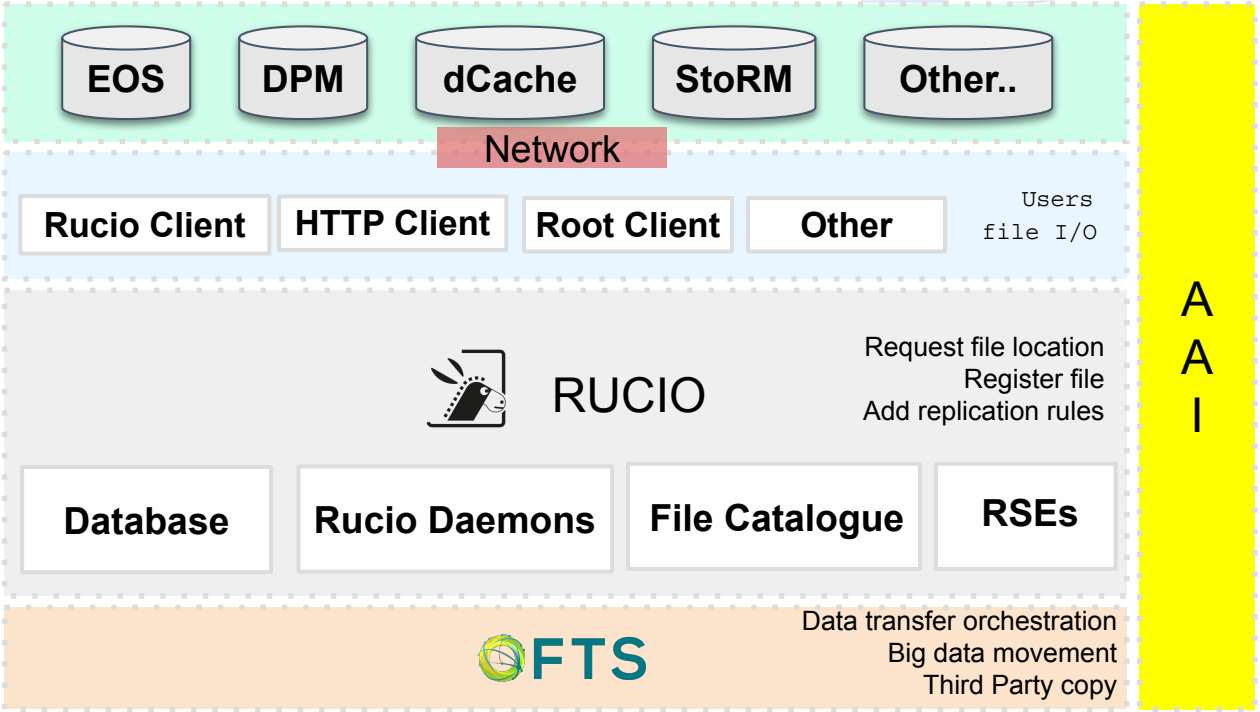



Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

ESCAPE
European Science Cluster of Astronomy & Particle physics ESFRI research Infrastructures

FTS

- Data Transfer orchestrator
- Can send Third Party Copy requests
- Compatible with most storage systems

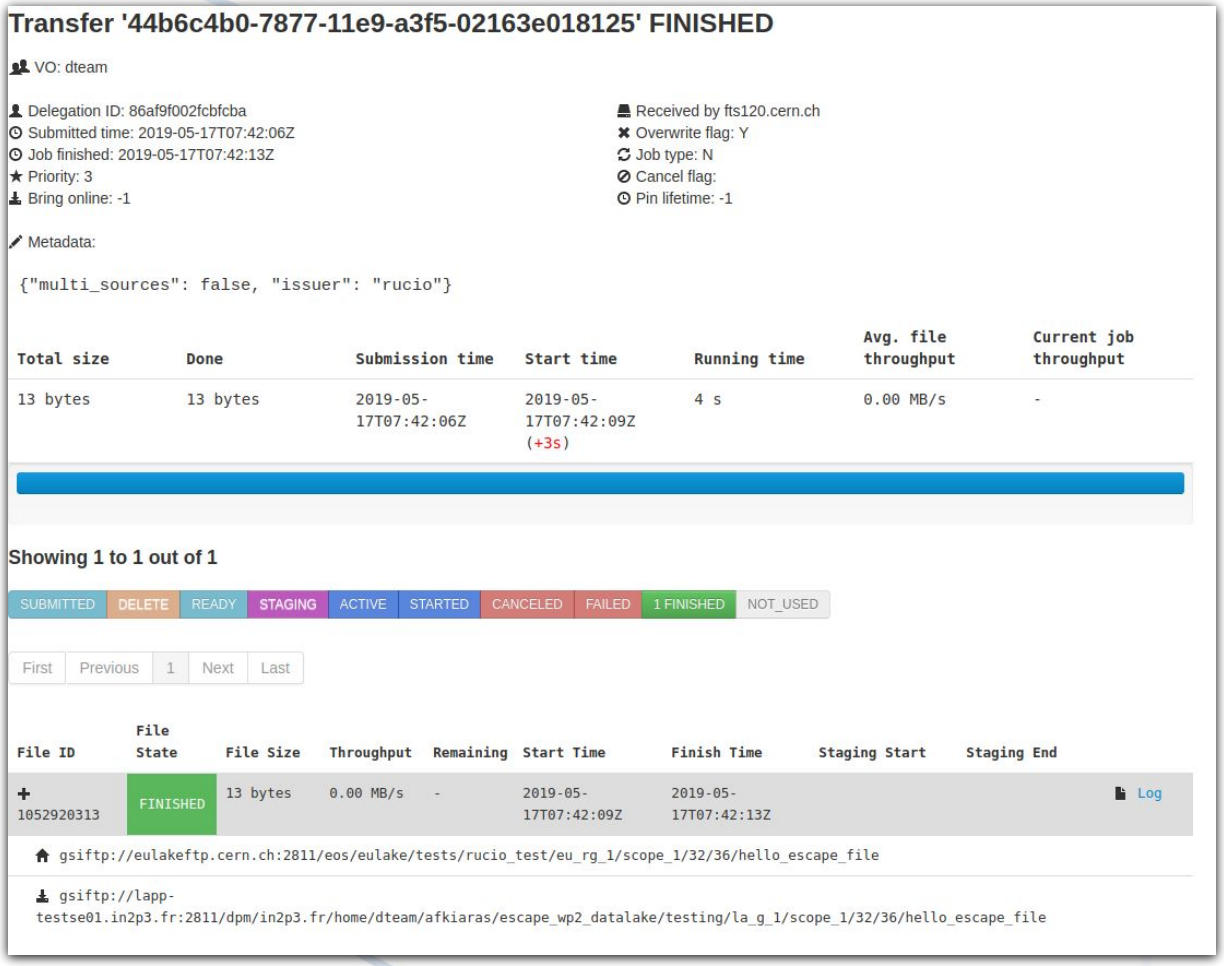

Transfer '44b6c4b0-7877-11e9-a3f5-02163e018125' FINISHED

VO: dteam

Delegation ID: 86af9f002fcbfcba
Submitted time: 2019-05-17T07:42:06Z
Job finished: 2019-05-17T07:42:13Z
Priority: 3
Bring online: -1

Received by fts120.cern.ch
Overwrite flag: Y
Job type: N
Cancel flag:
Pin lifetime: -1

Metadata:

```
{"multi_sources": false, "issuer": "rucio"}
```

| Total size | Done | Submission time | Start time | Running time | Avg. file throughput | Current job throughput |
|---|---|---|---|---|---|---|
| 13 bytes | 13 bytes | 2019-05-17T07:42:06Z | 2019-05-17T07:42:09Z (+3s) | 4 s | 0.00 MB/s | - |

Showing 1 to 1 out of 1

SUBMITTED | DELETE | READY | STAGING | ACTIVE | STARTED | CANCELED | FAILED | 1 FINISHED | NOT_USED

First | Previous | 1 | Next | Last

| File ID | File State | File Size | Throughput | Remaining | Start Time | Finish Time | Staging Start | Staging End | |
|---|---|---|---|---|---|---|---|---|---|
| 1052920313 | FINISHED | 13 bytes | 0.00 MB/s | - | 2019-05-17T07:42:09Z | 2019-05-17T07:42:13Z | | | Log |

gsiftp://eulakeftp.cern.ch:2811/eos/eulake/tests/rucio_test/eu_rg_1/scope_1/32/36/hello_escape_file

gsiftp://lapp-testse01.in2p3.fr:2811/dpm/in2p3.fr/home/dteam/afkiaras/escape_wp2_datalake/testing/la_g_1/scope_1/32/36/hello_escape_file


Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# Progress the last 4 months

- Set up ESCAPE specific Rucio Instance
- Connected EOS Pre-production Distributed Instance
- Connected LAPP DPM instance (Thanks Stephane, Frédérique, ...)
- Gained valuable experience with the Ecosystem
- Tried to document steps and lessons learned from setting this up


Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# Topology

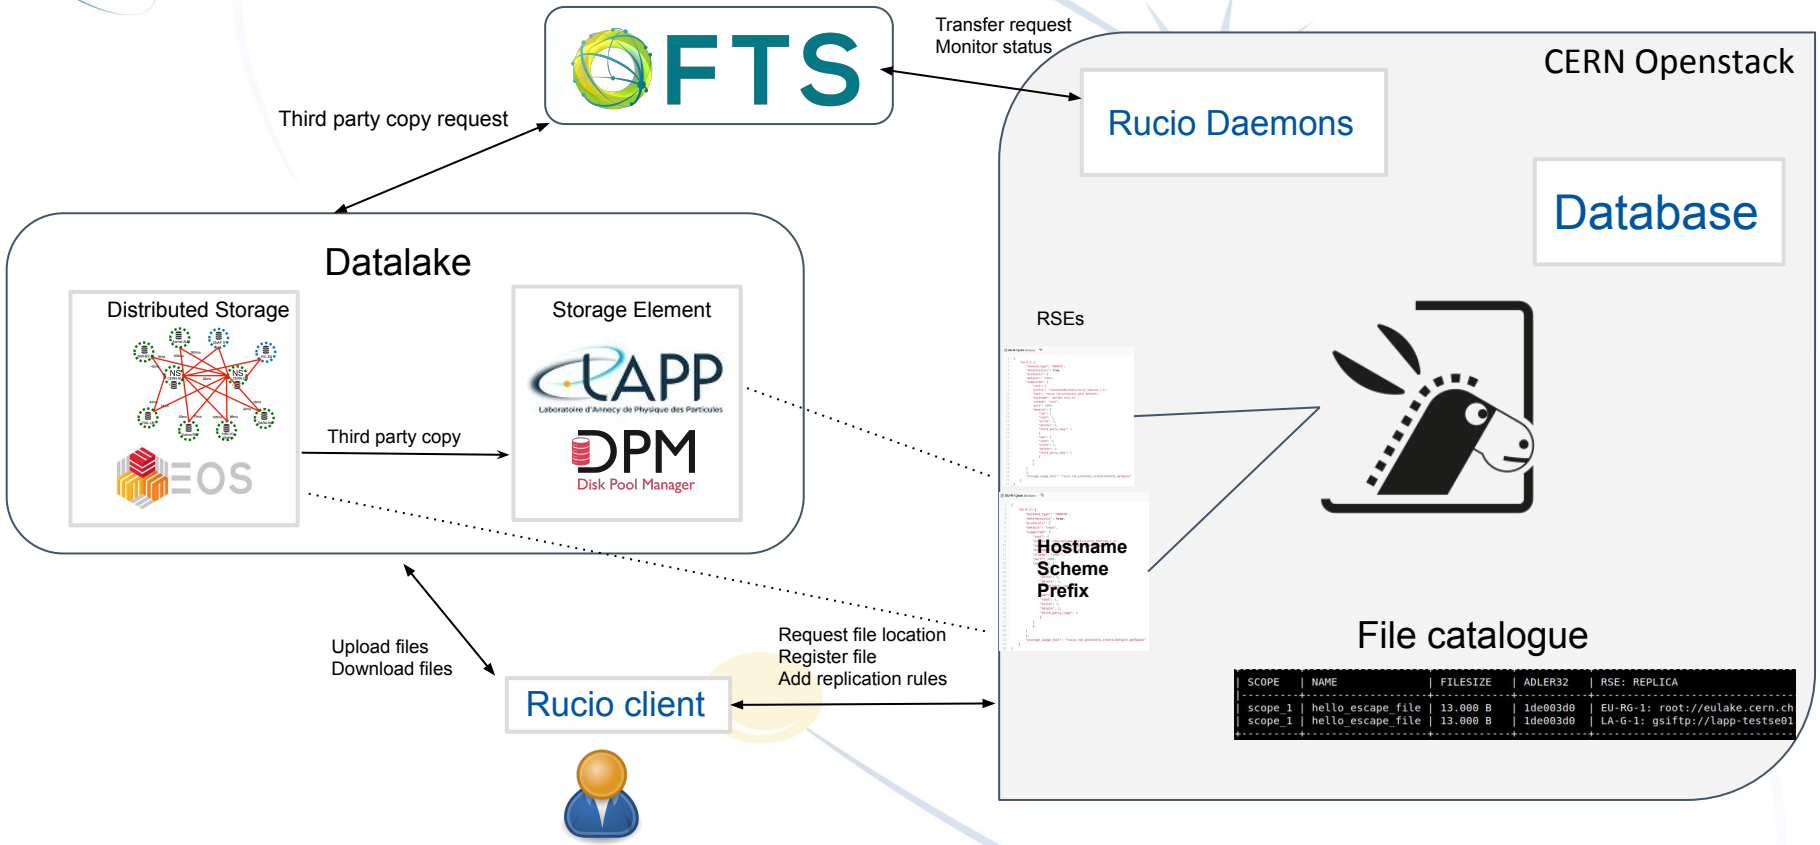



Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# Infrastructure

- Minimal Viable Infrastructure but can scale on Demand
- Use existing services but not lock in

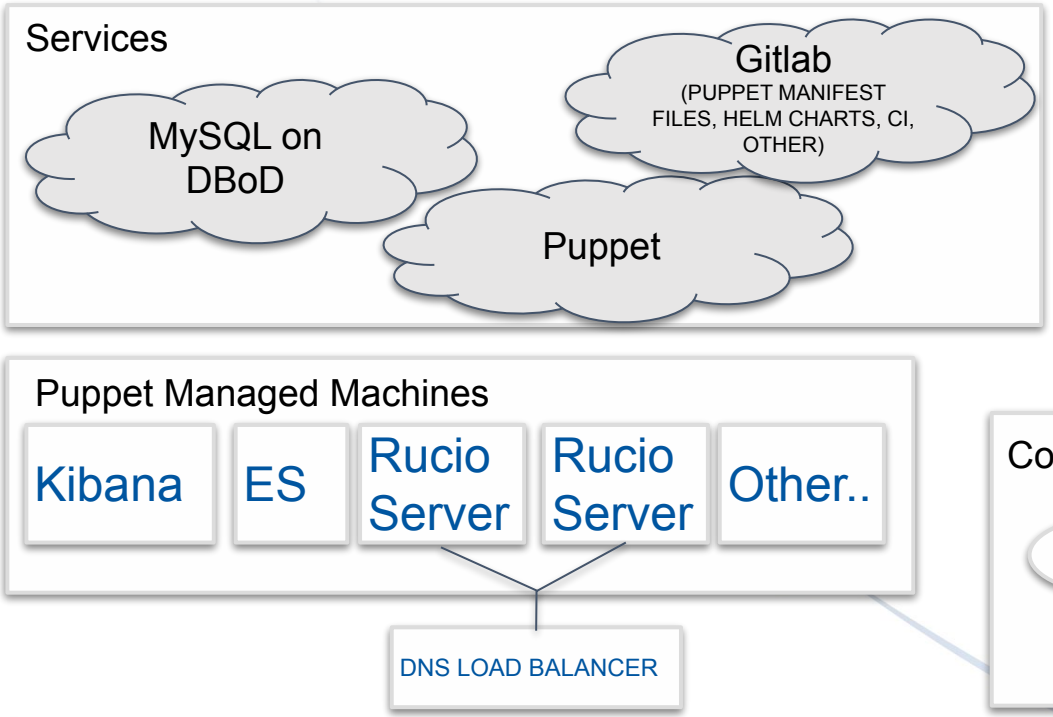



Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064

# Monitoring

- Internal Monitoring
  - Server Monitoring (CERN Monitoring Service stack, Collectd, InfluxDB, Grafana… )
  - Kubernetes Cluster Monitoring (Kubernetes Dashboard, Prometheus)
  - Rucio Server/Daemons Logs collection (Filebeat, logstash, Elasticsearch)
- Transfers Monitoring
  - ActiveMQ collecting Rucio Events (Hermes)
  - Storing “permanently” to Elasticsearch
  - Visualising with Kibana
- Access Monitoring
  - Collect file access traces from clients


Funded by the European Union’s Horizon 2020 - Grant N° 824064

# Documentation

Currently:

- Keeping notes and writing documentation on Google Docs, Google Keep, Jira Tickets, Gitbook, ESCAPE Notebook

Optimal:

- Have a single place for ESCAPE specific documentation only for the parts that is really necessary
- Contribute on the upstream docs where needed

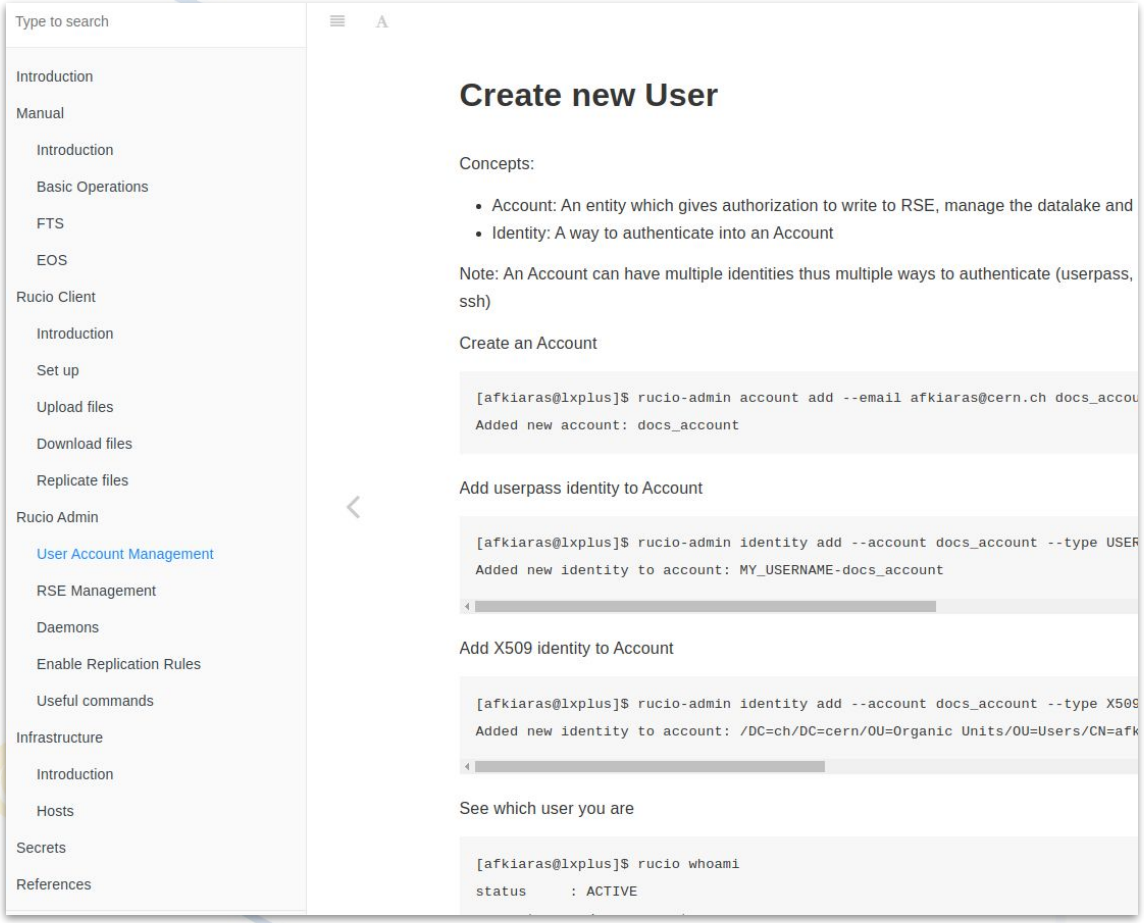



Funded by the European Union's Horizon 2020 - Grant N° 824064

# How to gain access / Try it out

- Rucio instance is currently accessible on the web
- Easiest authentication way is to handout userpass
- Since we don’t have AAI or an ESCAPE VO yet, its slightly complicated how your access propagates to the storage
  - For 1-2 people can manually add them to gridmap. For more we need an ESCAPE VO


Aristeidis.Fkiaras@cern.ch
Funded by the European Union’s Horizon 2020 - Grant N° 824064

# What this testbed is good enough for

- After connecting Storages (not production ones) provided by ESCAPE partners
  - Try experimental or even disruptive features
  - (Not everybody will be interested in the full rich features)

- Prototype or Test as they come out from the Rucio development team concepts that might be interesting to ESCAPE partners like:
  - QoS
  - AAI
  - Multi VO Rucio
  - Integration with CRIC ?
  - Event driven processing ?
  - Integration with Virtual Observatory ?
  - …..


Funded by the European Union's Horizon 2020 - Grant N° 824064

# Thank you!

Questions?

Comments!


Aristeidis.Fkiaras@cern.ch
Funded by the European Union's Horizon 2020 - Grant N° 824064