



ESCAPE

European Science Cluster of Astronomy &
Particle physics ESFRI research Infrastructures

ESCAPE prototypes a Data Infrastructure for Open Science

R. Bolton (SKAO), S. Campana (CERN), A. Ceccanti (INFN-CNAF),
X. Espinal (CERN), P. Fuhrmann (DESY), A. Fkiaras (CERN), Y. Grange
(ASTRON)

Computing in High Energy Physics, Adelaide, November 2019



ESCAPE

European Science Cluster of Astronomy & Particle physics
ESFRI research infrastructures

Horizon 2020
funded project




Goals:

Prototype an infrastructure adapted to the Exabyte-scale needs of the large science projects.

Ensure the sciences drive the development of the EOSC

Address *FAIR* data management



Data centres: CERN, INFN, DESY, GSI, Nikhef, SURFSara, RUG, CCIN2P3, PIC, LAPP, INAF

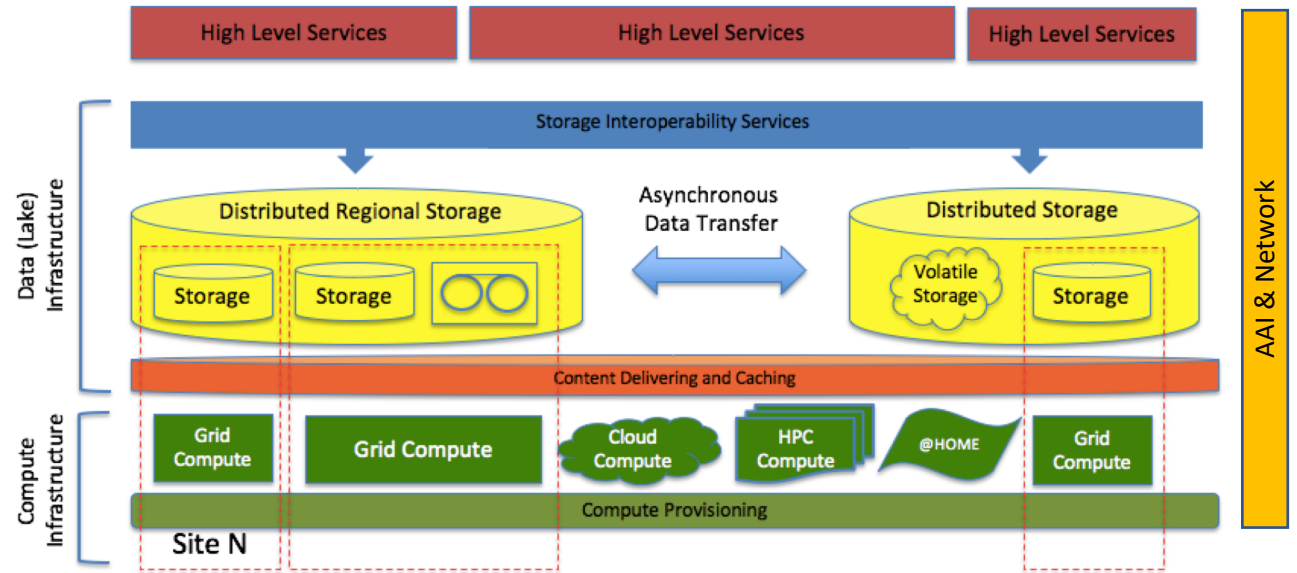
Science Projects

| | |
|-------------------|-------------------------|
| HL-LHC | SKA |
| FAIR | CTA |
| KM3Net | JIVE-ERIC |
| ELT | EST |
| EURO-VO (LSST) | EGO-VIRGO (CERN,ESO) |



WP2 – Data Infrastructure for Open Science

Data Lake building blocks



In synergy and complementing the work of WLCG DOMA: to define, integrate and commission an ecosystem of tools and services to build a data lake

Leaves to the science projects the flexibility to choose the services and layout most suitable to their needs. Provides a reference implementation

Contributes to deliver Open Access and FAIR data services: relies on trustable data repositories; enables data management policies; hides the complexities of the underlying infrastructure providing a transparent data access layer



WP2 important milestones

Today



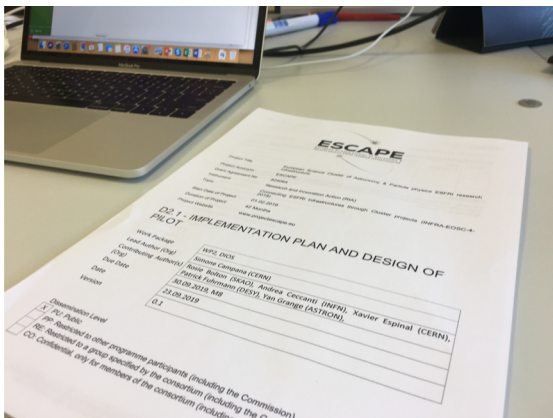
Implementation plan and the design of the of the pilot

Initial DataLake with at least 3 centers

Analysis and assessment of the first version of the pilot DataLake

Extended Prototype with more centres and tools

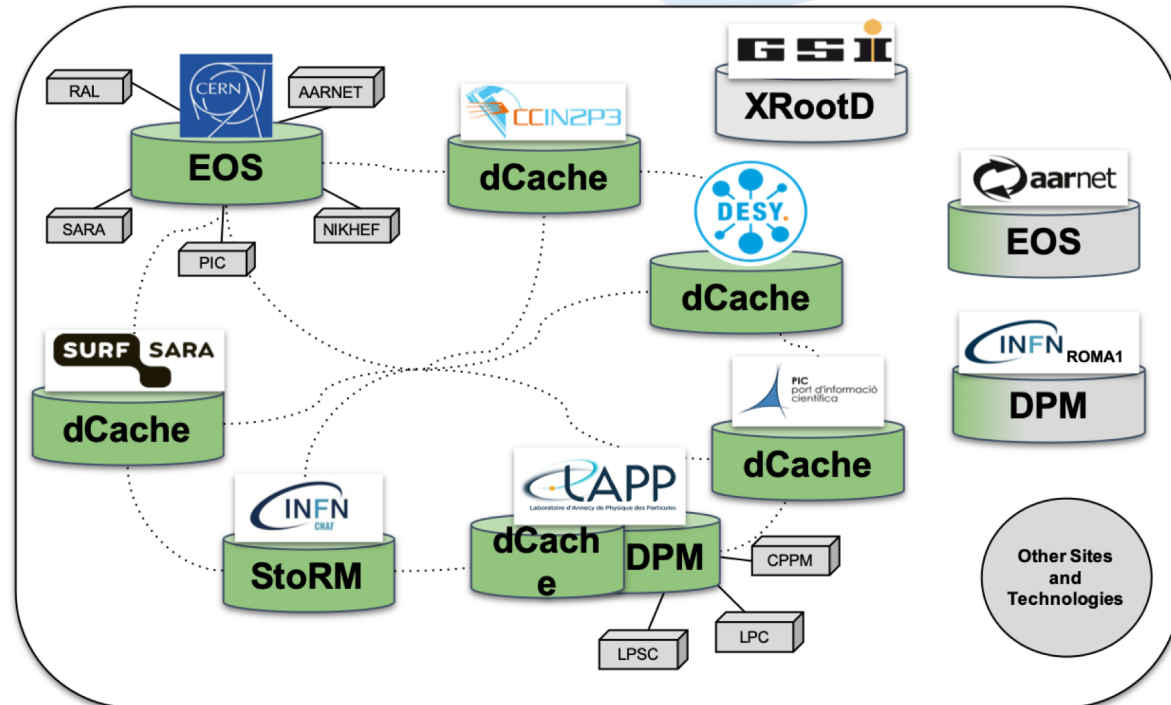
Final analysis and assessment of full prototype



Storage Technologies

EOS
DPM
dCache
StoRM
XRootD

http
xroot
gsiftp



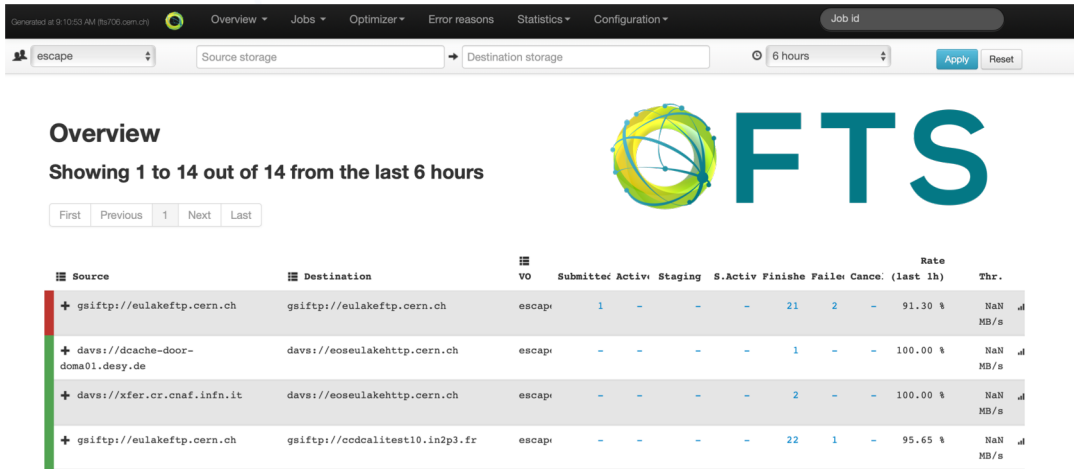
Currently deployed (**green**) and foreseen (**grey**) storage services in ESCAPE DataLake.

➤ Some centers joining even if not funded by ESCAPE!

An heterogeneous set of technologies, deployment models and Quality-Of-Service



Asynchronous Data Transfer



<https://fts3-pilot.cern.ch:8449/fts3/ftsmon/#/?vo=escape>


FTS is the workhorse for asynchronous point-to-point data transfer in the reference implementation

Planning for a gridFTP-free data lake: HTTP and xrootd

perfSONAR to provide low level network monitoring

Loss rate is <= 0.001% Loss rate is > 0.001% Loss rate is >= 0.1%

Found a total of 2 problems involving 2 hosts in the grid



ACHTUNG!

This is the WLCG OPN mesh. We should build an ESCAPE one



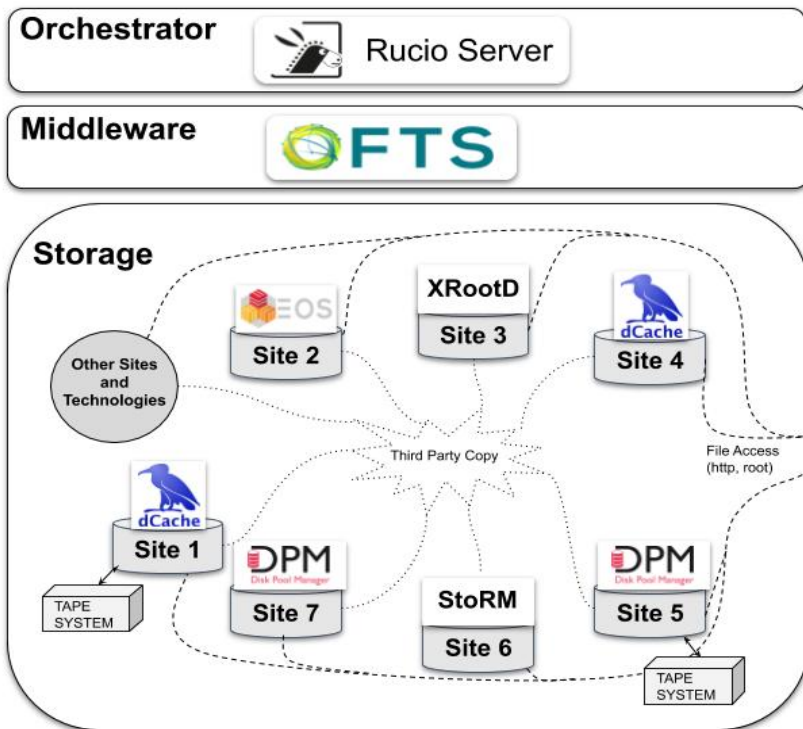
Orchestration Service



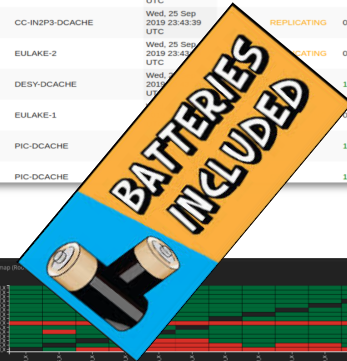
Rucio as the Orchestration service in the reference implementation

Rucio WebUI as interface for ESCAPE users

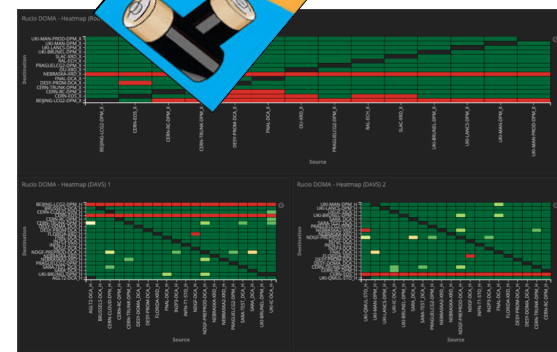
- File/dataset catalog, rule based engine




| Name | Account | RSE Expression | Creation Date | Remaining Lifetime | State | Locks OK | Locks Replicating | Locks Stuck |
|--|---------|-----------------|-------------------------------|--------------------|-------------|----------|-------------------|-------------|
| noise_auto_uploaded_xg0h1CNV0F03Aq09y1MQJH0tglJSMZ | root | SARA-DCACHE | Wed, 25 Sep 2019 23:43:40 UTC | - | STUCK | 0 | 0 | 1 |
| noise_auto_uploaded_xg0h1CNV0F03Aq09y1MQJH0tglJSMZ | root | CNAF-STORM | Wed, 25 Sep 2019 23:43:40 UTC | - | OK | 1 | 0 | 0 |
| noise_auto_uploaded_xg0h1CNV0F03Aq09y1MQJH0tglJSMZ | root | CC-IN2P3-OCACHE | Wed, 25 Sep 2019 23:43:39 UTC | - | REPLICATING | 0 | 1 | 0 |
| noise_auto_uploaded_xg0h1CNV0F03Aq09y1MQJH0tglJSMZ | root | EULAKE-2 | Wed, 25 Sep 2019 23:43:39 UTC | - | PENDING | 0 | 1 | 0 |
| noise_auto_uploaded_xg0h1CNV0F03Aq09y1MQJH0tglJSMZ | root | DESY-DCACHE | Wed, 25 Sep 2019 23:43:39 UTC | - | OK | 1 | 0 | 0 |
| noise_auto_uploaded_xg0h1CNV0F03Aq09y1MQJH0tglJSMZ | root | EULAKE-1 | Wed, 25 Sep 2019 23:43:39 UTC | - | OK | 1 | 0 | 0 |
| noise_auto_uploaded_xg0h1CNV0F03Aq09y1MQJH0tglJSMZ | root | PIC-DCACHE | Wed, 25 Sep 2019 23:43:39 UTC | - | OK | 1 | 0 | 0 |
| noise_auto_uploaded_g0Kya1bKhrjKb8HSpkQZYarB8Zn1 | root | PIC-DCACHE | Wed, 25 Sep 2019 23:43:39 UTC | - | OK | 1 | 0 | 0 |



Functional Tests & Visualization

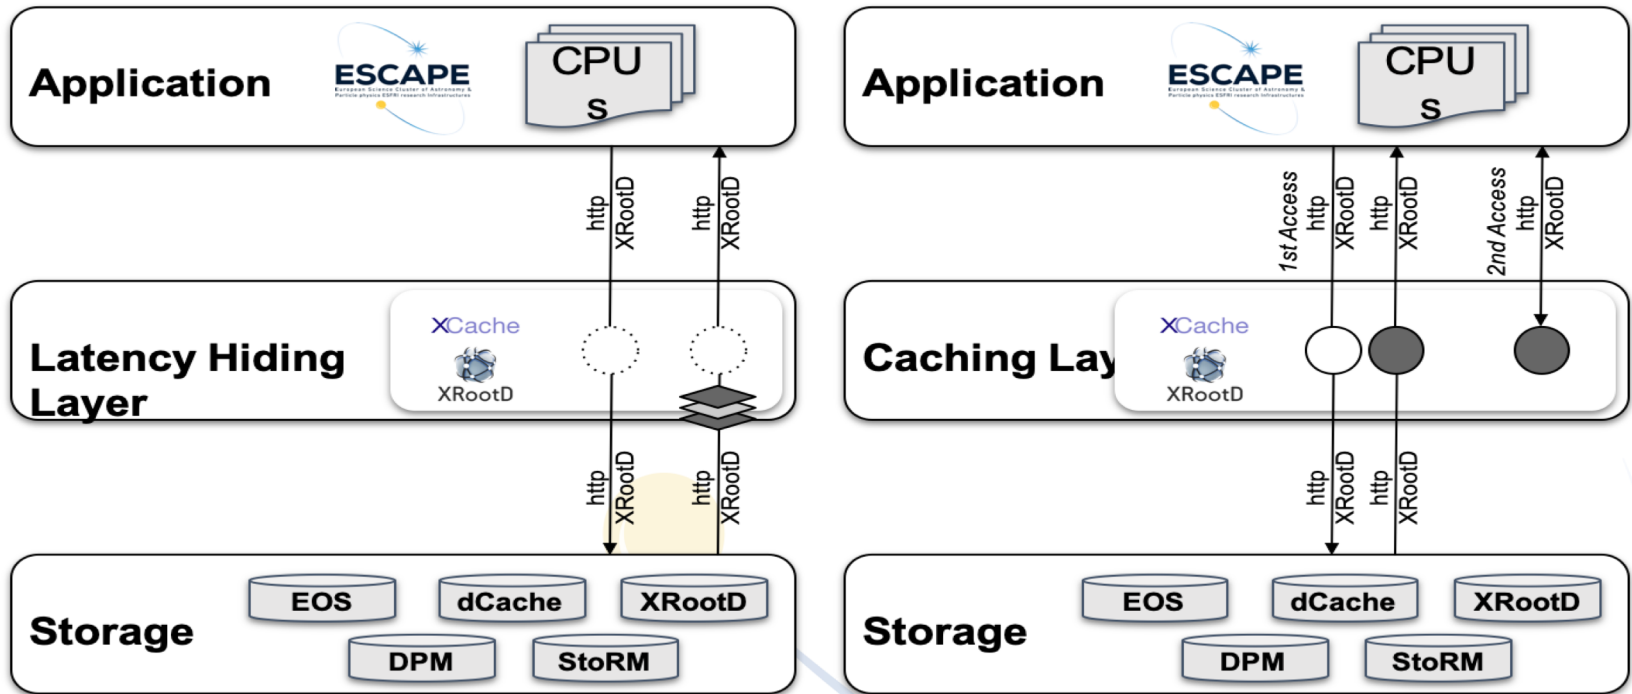


Caching and Latency hiding



xCache technology as reference implementation.

Support HTTP and xrootd protocols



Authentication, Authorization, Identity

AAI has a much broader scope than WP2

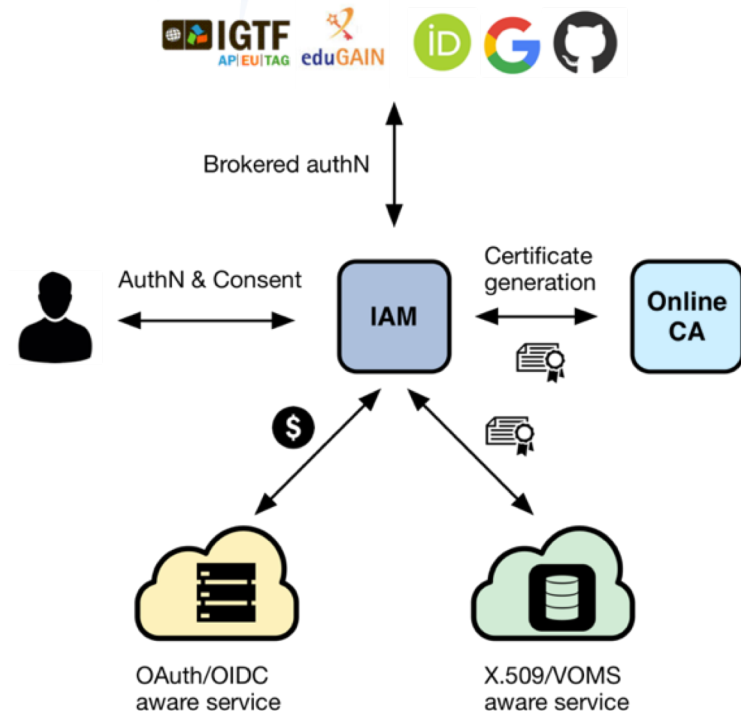
IAM was chosen as reference implementation service for AAI

Token-based authentication rather than X509 as baseline. IAM enables translation X509-to-Token for coexistence period

X509-free DataLake is not realistic in the timescale of ESCAPE (IMHO)

Some fully token based use cases however are in range

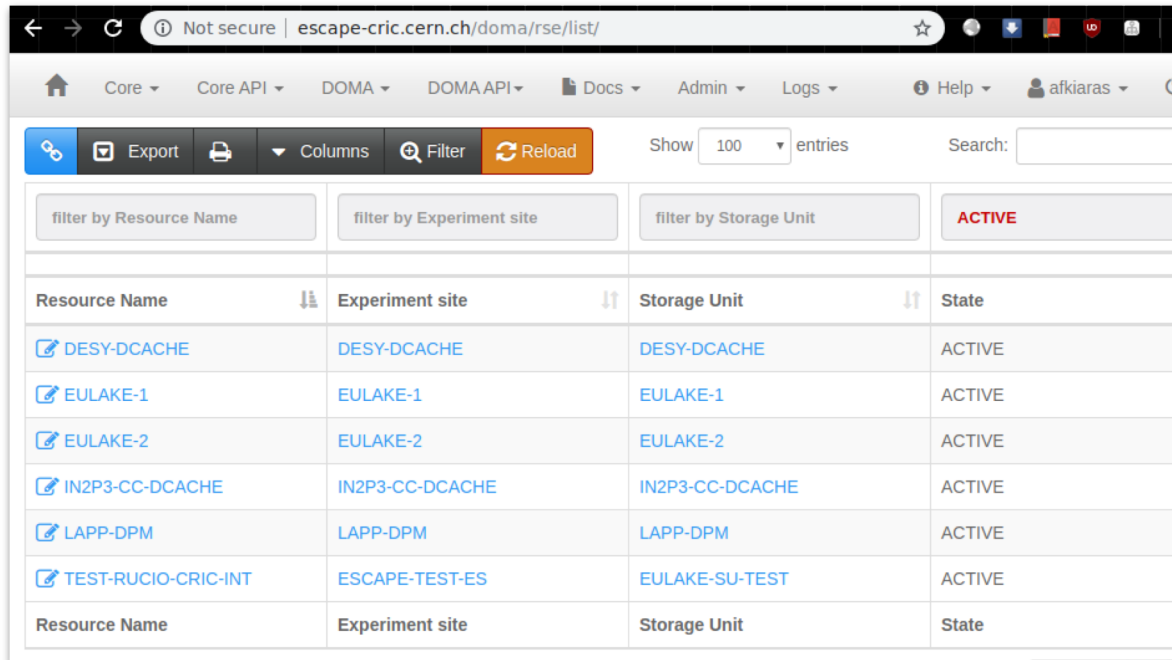
Leveraging the work in the Federated Identity Management For Research (FIM4R) initiative



Configuration

CRIC as the DataLake configuration service

- In use in WLCG, integrated with Rucio

Not secure | escape-cric.cern.ch/doma/rse/list/

Core Core API DOMA DOMA API Docs Admin Logs Help afkiaras

Export Columns Filter Reload Show 100 entries Search:

filter by Resource Name filter by Experiment site filter by Storage Unit ACTIVE

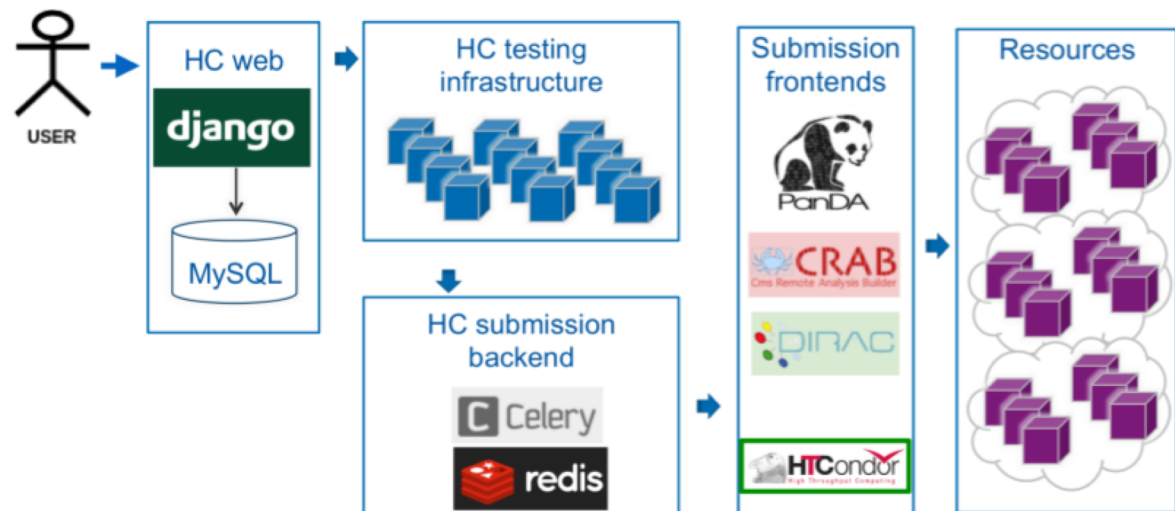
| Resource Name | Experiment site | Storage Unit | State |
|-------------------------------------|-----------------|-----------------|--------|
| DESY-DCACHE | DESY-DCACHE | DESY-DCACHE | ACTIVE |
| EULAKE-1 | EULAKE-1 | EULAKE-1 | ACTIVE |
| EULAKE-2 | EULAKE-2 | EULAKE-2 | ACTIVE |
| IN2P3-CC-DCACHE | IN2P3-CC-DCACHE | IN2P3-CC-DCACHE | ACTIVE |
| LAPP-DPM | LAPP-DPM | LAPP-DPM | ACTIVE |
| TEST-RUCIO-CRIC-INT | ESCAPE-TEST-ES | EULAKE-SU-TEST | ACTIVE |
| Resource Name | Experiment site | Storage Unit | State |



Commissioning

Hammercloud for testing and commissioning

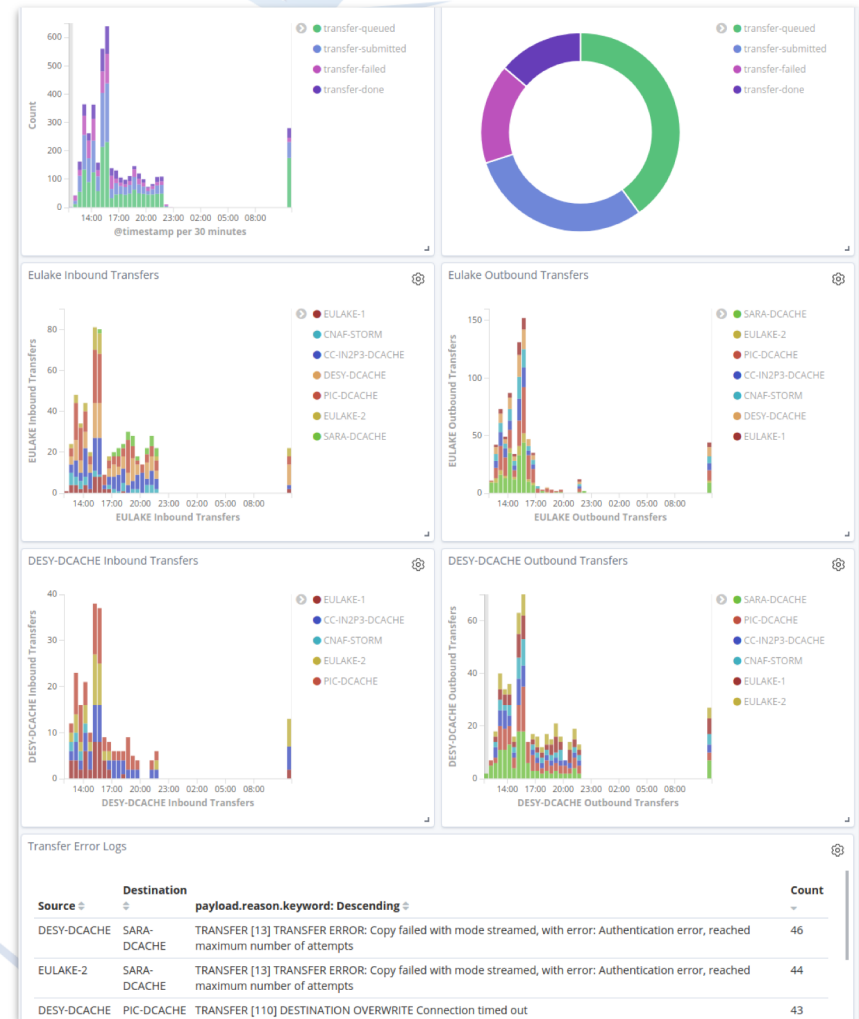
- Can run applications in distributed environments, through different plugins
- Comes with monitoring and analytics



Monitoring

Plan to build an ESCAPE DataLake dashboard, based on open source tools

Most of the information (Rucio, FTS) are already collected in ElasticSearch and can be exposed e.g. via Kibana



Conclusions



As part of the ESCAPE project we are contributing to prototype a Data Infrastructure for the large volume of Open Science data in the 2020s

Building on the collaboration between High Energy and Nuclear Physics and Astronomy, driven by the sciences.

In synergy and leveraging existing projects and solutions