



# ESCAPE

European Science Cluster of Astronomy &  
Particle physics ESFRI research Infrastructures

## WP2 Fortnightly Meeting - Datalake Status

Aris Fkiaras

CERN

October 9th, 2019



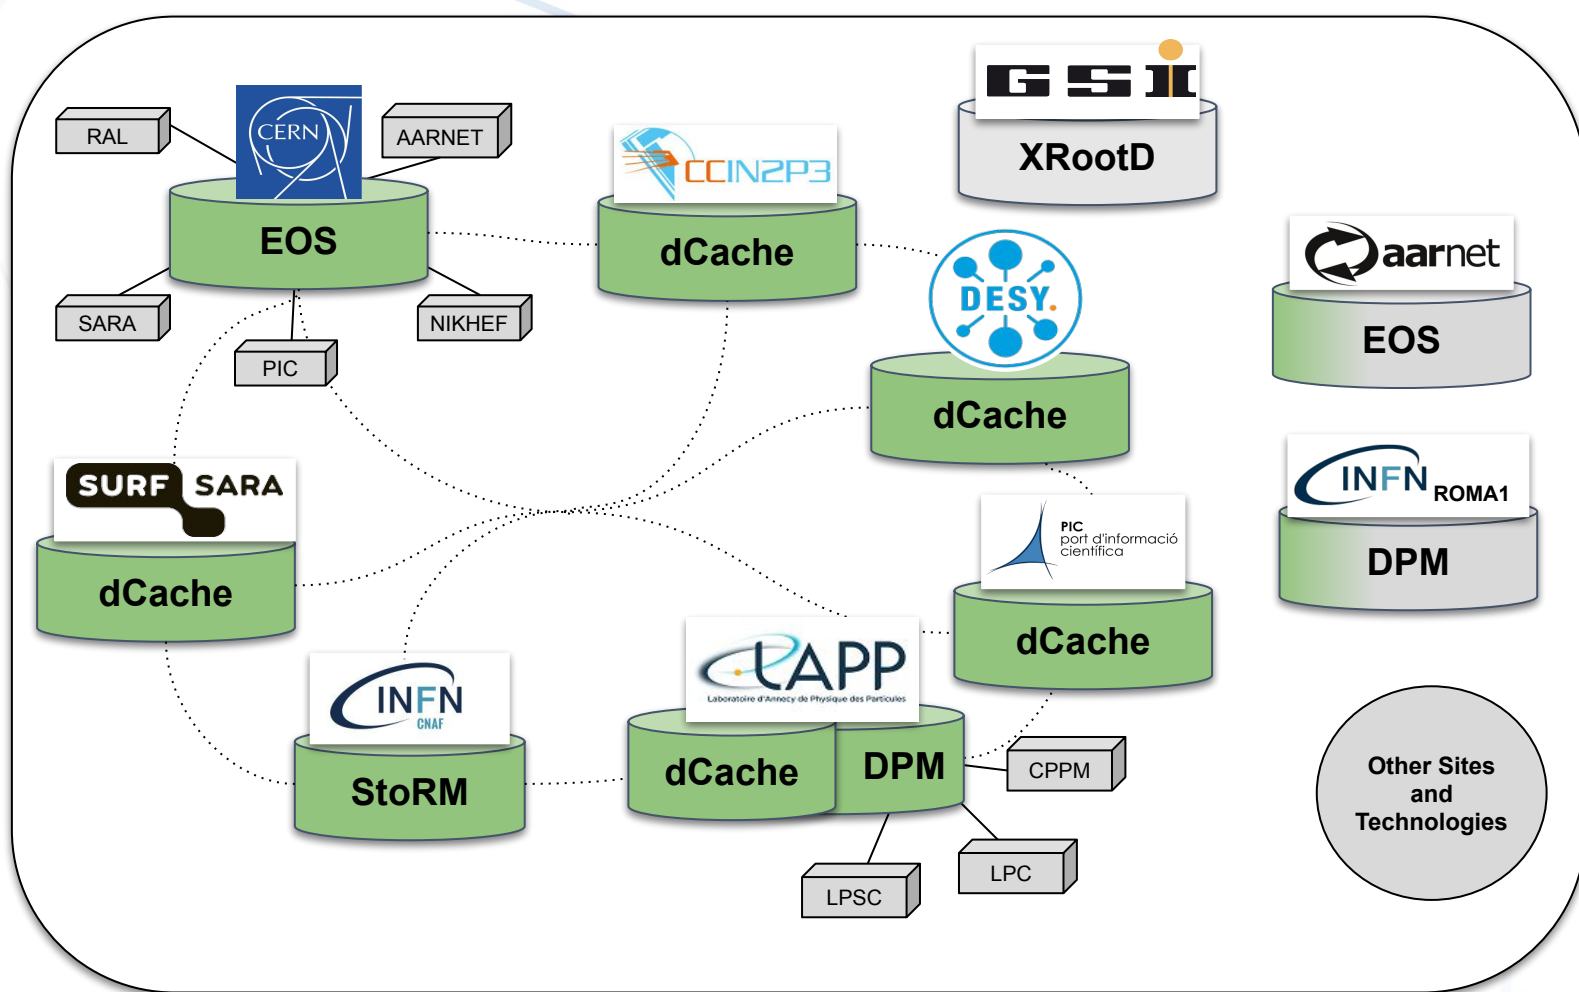
# Overview

- Datalake Status
- Where to find latest information
- Ways to access the datalake
- Monitoring solutions in place and feature plans



EOS ☒  
DPM ☒  
dCache ☒  
StoRM ☒  
XRootD ☐

http ☒  
xroot ☒  
gsiftp ☒



# “Realtime” detailed Status

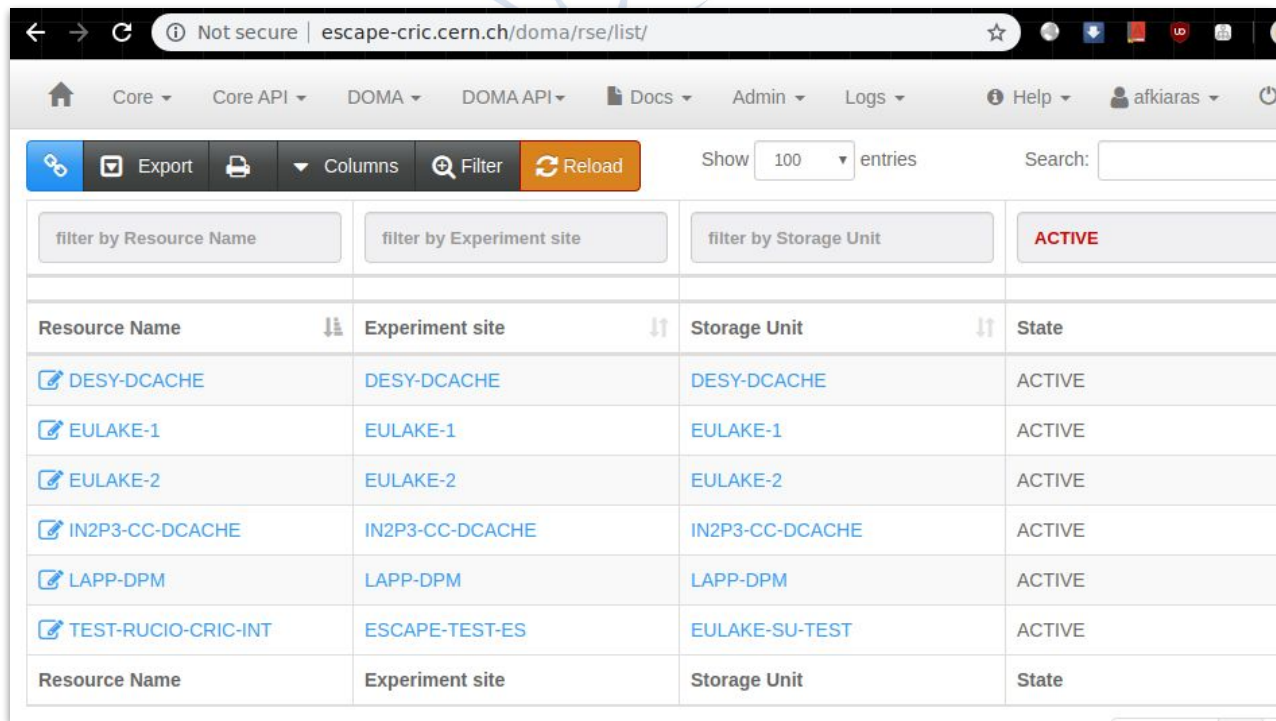
- Available on the official ESCAPE wiki  
[https://wiki.escape2020.de/index.php/WP2 - DIOS - Datalake Status](https://wiki.escape2020.de/index.php/WP2_-_DIOS_-_Datalake_Status)
- Please update missing or invalid information

Institute	Technology	Version	Quota	XROOT	HTTPS	GSIFT	ESCAPE VO enabled (X509)	Token based auth/z
CERN	EOS	?	?	YES	YES	YES	WiP	NO
LAPP	DPM	?	?	YES	YES	YES	WiP	?
SURFsara	dCache	?	?	YES	YES	YES	YES	?
LAPP	dCache	?	?	NO	YES (http)	NO	YES	?
GSI	xRootD	?	?	?	?	?	?	?
INFN-CNAF	StoRM	?	?	NO	YES	NO	YES	?
INFN-ROMA1	DPM	?	?	WiP	WiP	WiP	WiP	?
DESY	dCache	?	?	YES	YES	YES	YES	?
CC-IN2P3	dCache	?	60 TB	YES	YES	YES	YES	?



# CRIC: Information System

- Started populating with ESCAPE storage details
- At the moment stores only static information (protocols, endpoints, ports)
- Will work on the synchronization with Rucio during the Rucio Coding Camp (15-17 Oct)



The screenshot shows a web browser at the URL `escape-cric.cern.ch/doma/rse/list/`. The interface includes a navigation bar with links for Core, Core API, DOMA, DOMA API, Docs, Admin, Logs, Help, and a user profile for 'afkiaras'. Below the navigation bar are buttons for 'Export', 'Columns', 'Filter', and 'Reload'. A table displays a list of storage resources with columns for Resource Name, Experiment site, Storage Unit, and State. The table is filtered to show only 'ACTIVE' resources. The data rows are as follows:

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



# Rucio client

- Accessing the datalake using the Rucio client

```
(rucio) [afkiaras@lxplus761 postgres]$ rucio upload hello_wp2_meeting --rse PIC-DCACHE --scope testing
2019-10-09 08:59:44,125 INFO      Preparing upload for file hello_wp2_meeting
2019-10-09 08:59:44,145 INFO      File DID already exists
2019-10-09 08:59:44,330 INFO      Successfully added replica in Rucio catalogue at PIC-DCACHE
2019-10-09 08:59:44,596 INFO      Trying upload with davs to PIC-DCACHE
2019-10-09 08:59:45,593 INFO      Successfully uploaded file hello_wp2_meeting
```

- Instructions on how to set up the Rucio client

[https://rucio.readthedocs.io/en/latest/installing\\_clients.html](https://rucio.readthedocs.io/en/latest/installing_clients.html)

- Features:

- Upload/Download files
- Add Rucio Rules (Then Rucio triggers FTS transfers to satisfy those rules)
- Configure the datalake (rucio-admin client)
- ....





- Enabled features:
  - Monitor transfer status
  - Add new Rules
  - Download files
  - Monitor user quota per RSE

Rucio UI Monitoring Data Transfers (R2D2) pattern OR name OR rule id Search Using account: root

You are here: Rucio Rule Definition Droid - List Rules Rucio Version ( WebUI / Server ): 1.20.3 / 1.20.6

Rules 1 New

Account RSE State Activity Interval

root RSE User Subscriptions 14 days Apply

Show 100 entries Search:

Name	Account	RSE Expression	Creation Date	Remaining Lifetime	State	Locks OK	Locks Replicating	Locks Stuck
<a href="#">noise:auto_uploaded_Xq0h1CfV0F03AqhO9y1MQJH0t gLjSM8Z</a>	root	SARA-DCACHE	Wed, 25 Sep 2019 23:43:40 UTC	-	STUCK	0	0	1
<a href="#">noise:auto_uploaded_Xq0h1CfV0F03AqhO9y1MQJH0t gLjSM8Z</a>	root	CNAF-STORM	Wed, 25 Sep 2019 23:43:40 UTC	-	OK	1	0	0
<a href="#">noise:auto_uploaded_Xq0h1CfV0F03AqhO9y1MQJH0t gLjSM8Z</a>	root	CC-IN2P3-DCACHE	Wed, 25 Sep 2019 23:43:39 UTC	-	REPLICATING	0	1	0
<a href="#">noise:auto_uploaded_Xq0h1CfV0F03AqhO9y1MQJH0t gLjSM8Z</a>	root	EULAKE-2	Wed, 25 Sep 2019 23:43:38 UTC	-	REPLICATING	0	1	0
<a href="#">noise:auto_uploaded_Xq0h1CfV0F03AqhO9y1MQJH0t gLjSM8Z</a>	root	DESY-DCACHE	Wed, 25 Sep 2019 23:43:38 UTC	-	OK	1	0	0
<a href="#">noise:auto_uploaded_Xq0h1CfV0F03AqhO9y1MQJH0t gLjSM8Z</a>	root	EULAKE-1	Wed, 25 Sep 2019 23:43:37 UTC	-	REPLICATING	0	1	0
<a href="#">noise:auto_uploaded_Xq0h1CfV0F03AqhO9y1MQJH0t gLjSM8Z</a>	root	PIC-DCACHE	Wed, 25 Sep 2019 23:43:35 UTC	-	OK	1	0	0
<a href="#">noise:auto_uploaded_g0Kya1bKhrijkbBrisjxQZYarb8 R201</a>	root	PIC-DCACHE	Wed, 25 Sep 2019 23:43:34 UTC	-	OK	1	0	0



# Functional tests of the Datalake

- Cron job running every half hour
- Uses the Rucio Client
- For each SITE in the datalake:
  - Creates a unique file (10MB atm)
  - Uploads the file using the Rucio Client
    - File is uploaded to the site
    - File is registered on Rucio catalogue
  - Adds Rucio rules requesting that file to exist on all other sites as well
  - Rucio daemons notice the Rules and send requests to FTS
  - FTS executes transfers using Third Party Copy
  - Rucio periodically checks FTS and updates catalogue accordingly





# FTS Dashboard

- Realtime monitoring
- Maximum 6 hours of transfers
- Can get detailed transfer logs

[https://fts3-pilot.cern.ch:8449/fts3/ftsmon/#/?vo=escape&source\\_se=&dest\\_se=&time\\_window=6](https://fts3-pilot.cern.ch:8449/fts3/ftsmon/#/?vo=escape&source_se=&dest_se=&time_window=6)

## Overview

Showing 1 to 44 out of 44 from the last 6 hours

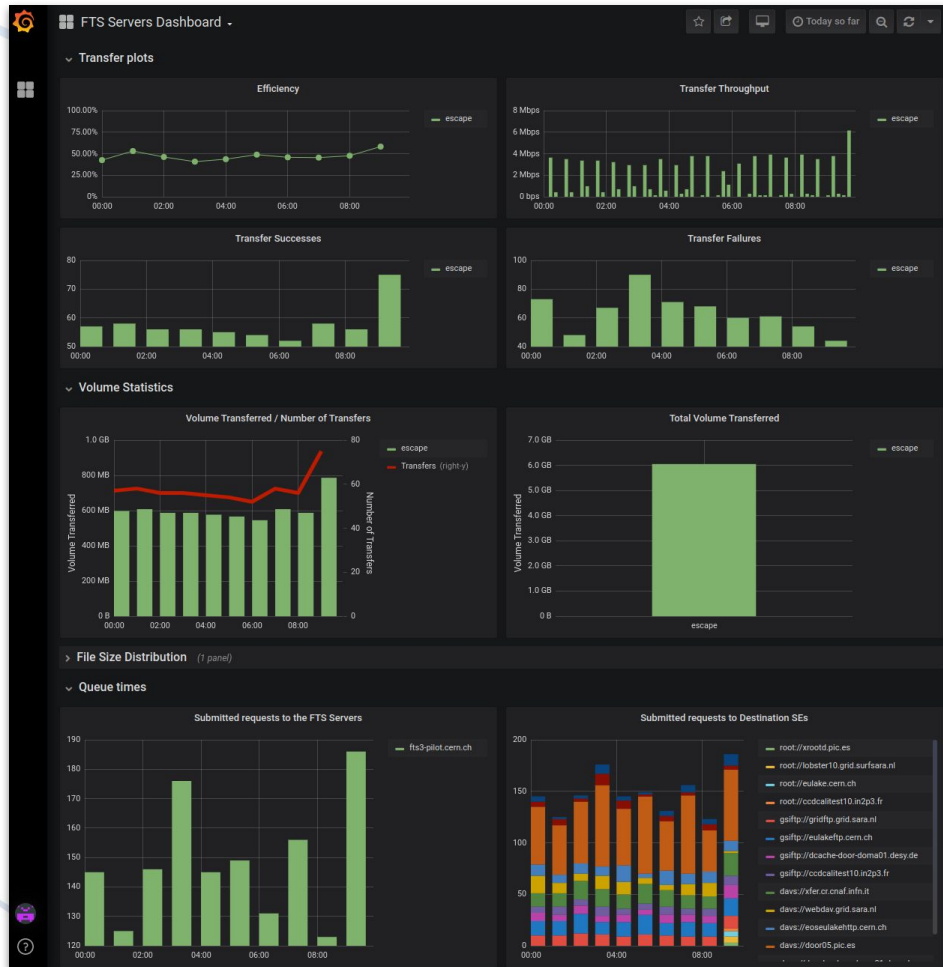
First Previous 1 Next Last

Source	Destination	VO	Submitte	Activ	Staging	S.Activ	Finish	Fail	Cance	Rate (last 1h)	Thr.
+ root://ccdcalttest10.in2p3.fr	root://lobster10.grid.surfsara.nl	escap	-	2	-	-	-	-	-	-	-
+ root://eulake.cern.ch	root://lobster10.grid.surfsara.nl	escap	-	1	-	-	-	-	-	-	-
+ davs://door05.pic.es	davs://xfer.cr.cnaf.infn.it	escap	-	1	-	-	3	15	-	16.67 %	-
+ davs://dcache-door-doma01.desy.de	davs://door05.pic.es	escap	-	1	-	-	5	43	-	10.42 %	NaN MB/s
+ davs://door05.pic.es	davs://dcache-door-doma01.desy.de	escap	-	1	-	-	5	14	-	26.32 %	-
+ root://xrootd.pic.es	root://lobster10.grid.surfsara.nl	escap	-	1	-	-	-	-	-	-	-
+ davs://eoseulakehttp.cern.ch	davs://xfer.cr.cnaf.infn.it	escap	-	1	-	-	28	1	-	96.55 %	0.03 MB/s
+ root://eulake.cern.ch	root://ccdcalttest10.in2p3.fr	escap	-	-	-	-	-	2	-	0.00 %	NaN MB/s
+ davs://dcache-door-doma01.desy.de	davs://eoseulakehttp.cern.ch	escap	-	-	-	-	2	-	-	100.00 %	NaN MB/s
+ davs://dcache-door-doma01.desy.de	davs://webdav.grid.sara.nl	escap	-	-	-	-	-	7	-	0.00 %	NaN MB/s
+ davs://eoseulakehttp.cern.ch	davs://webdav.grid.sara.nl	escap	-	-	-	-	-	8	-	0.00 %	NaN MB/s
+ root://xrootd.pic.es	root://ccdcalttest10.in2p3.fr	escap	-	-	-	-	-	1	-	0.00 %	NaN MB/s
+ davs://dcache-door-doma01.desy.de	davs://ccdcalttest10.in2p3.fr	escap	-	-	-	-	1	-	-	100.00 %	NaN MB/s
+ gsiftp://dcache-door-doma01.desy.de	gsiftp://ccdcalttest10.in2p3.fr	escap	-	-	-	-	15	-	-	100.00 %	NaN MB/s
+ root://eulake.cern.ch	root://eulake.cern.ch	escap	-	-	-	-	2	-	-	100.00 %	NaN MB/s
+ davs://eoseulakehttp.cern.ch	davs://door05.pic.es	escap	-	-	-	-	18	77	-	18.95 %	0.01 MB/s
+ gsiftp://eulakehttp.cern.ch	gsiftp://ccdcalttest10.in2p3.fr	escap	-	-	-	-	26	-	-	100.00 %	0.07 MB/s
+ davs://eoseulakehttp.cern.ch	davs://dcache-door-doma01.desy.de	escap	-	-	-	-	1	-	-	100.00 %	NaN MB/s
+ davs://eoseulakehttp.cern.ch	davs://ccdcalttest10.in2p3.fr	escap	-	-	-	-	2	-	-	100.00 %	NaN MB/s



# FTS Monitoring

- Official FTS Monitoring dashboards on CERN-Monit
- Stores information long term
- Managed by FTS team thus can not customize it a lot for ESCAPE needs



# Custom solution on Elastic Stack

- Collects information coming from Rucio
  - FTS transfer status events
  - Requests by users like dataset deletion
  - Server logs
- Customizable since it is managed by us
- Not that stable or performant



# Monitoring plans

- Implement a matrix showing the results of the functional tests
  - Involving all partner sites
    - All storage flavours represented (dCache, DPM, STORM, EOS and xrootd)
    - Mixed protocol scenario (xrootd, http and gsiftp) but http enabled everywhere
- Collect information of other datalake components
  - Xcache, rucio client traces, rucio events
- Visualise everything in a common Monitoring Dashboard
  - Grafana on CERN-Monit?



# Thank you!

Comments?  
Questions?



# Backup slides





# FTS zoom

