



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

# WP2 24h stress test

Yan Grange

Xavier Espinal, Rosie Bolton, Riccaro di Maria, Paul Millar, Andrea Ceccanti

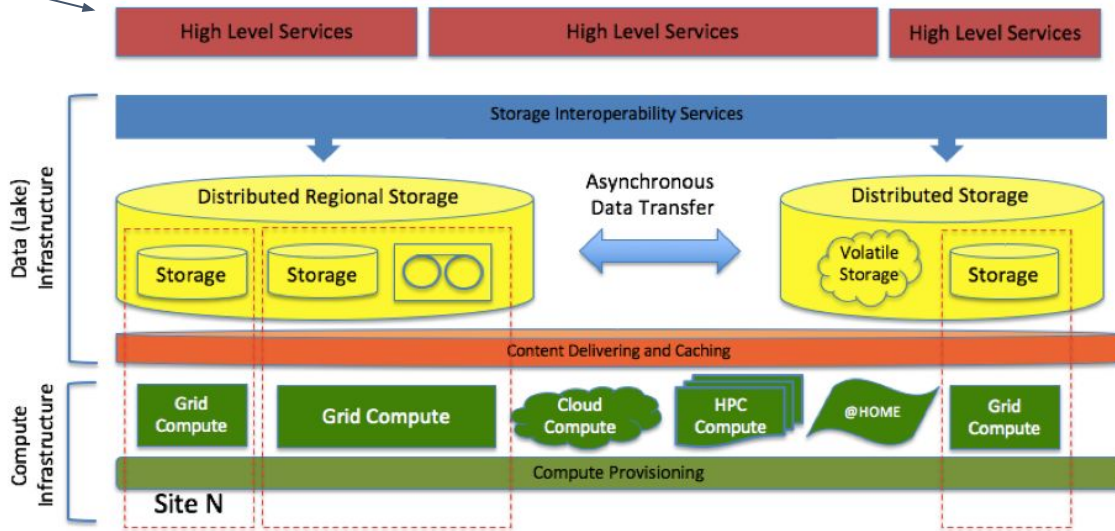
October 5, 2020 - WP5 tech talk



You are here!

# WP2 refresher!

Data Lake building blocks



Define, integrate and commission an ecosystem of tools and services to build a data lake

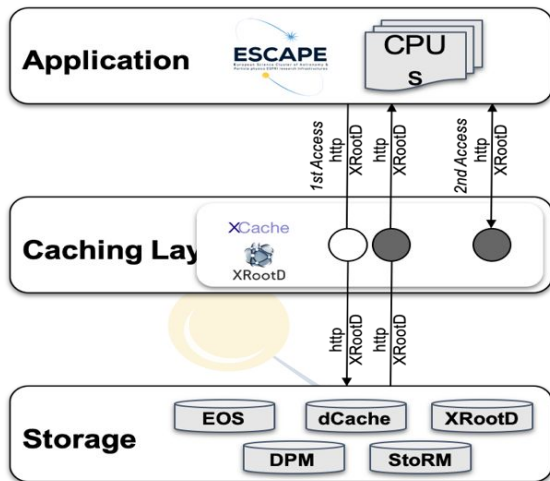
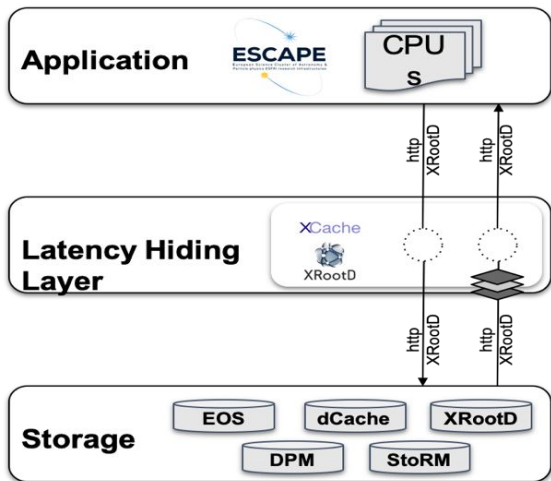
Leaves to the science projects the flexibility to choose the services and layout most suitable to their needs. Provides a reference implementation

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

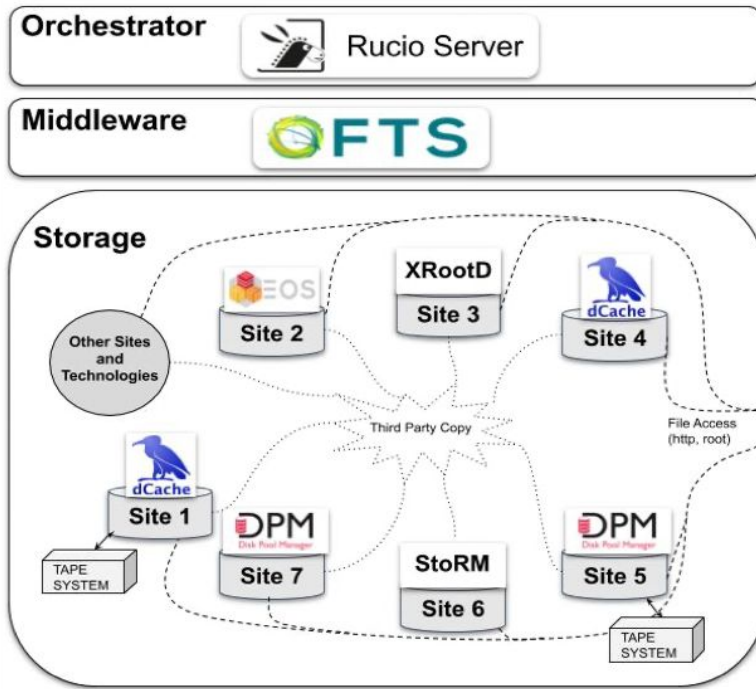


# WP2 refresher (2)!

## Caching






## Data management



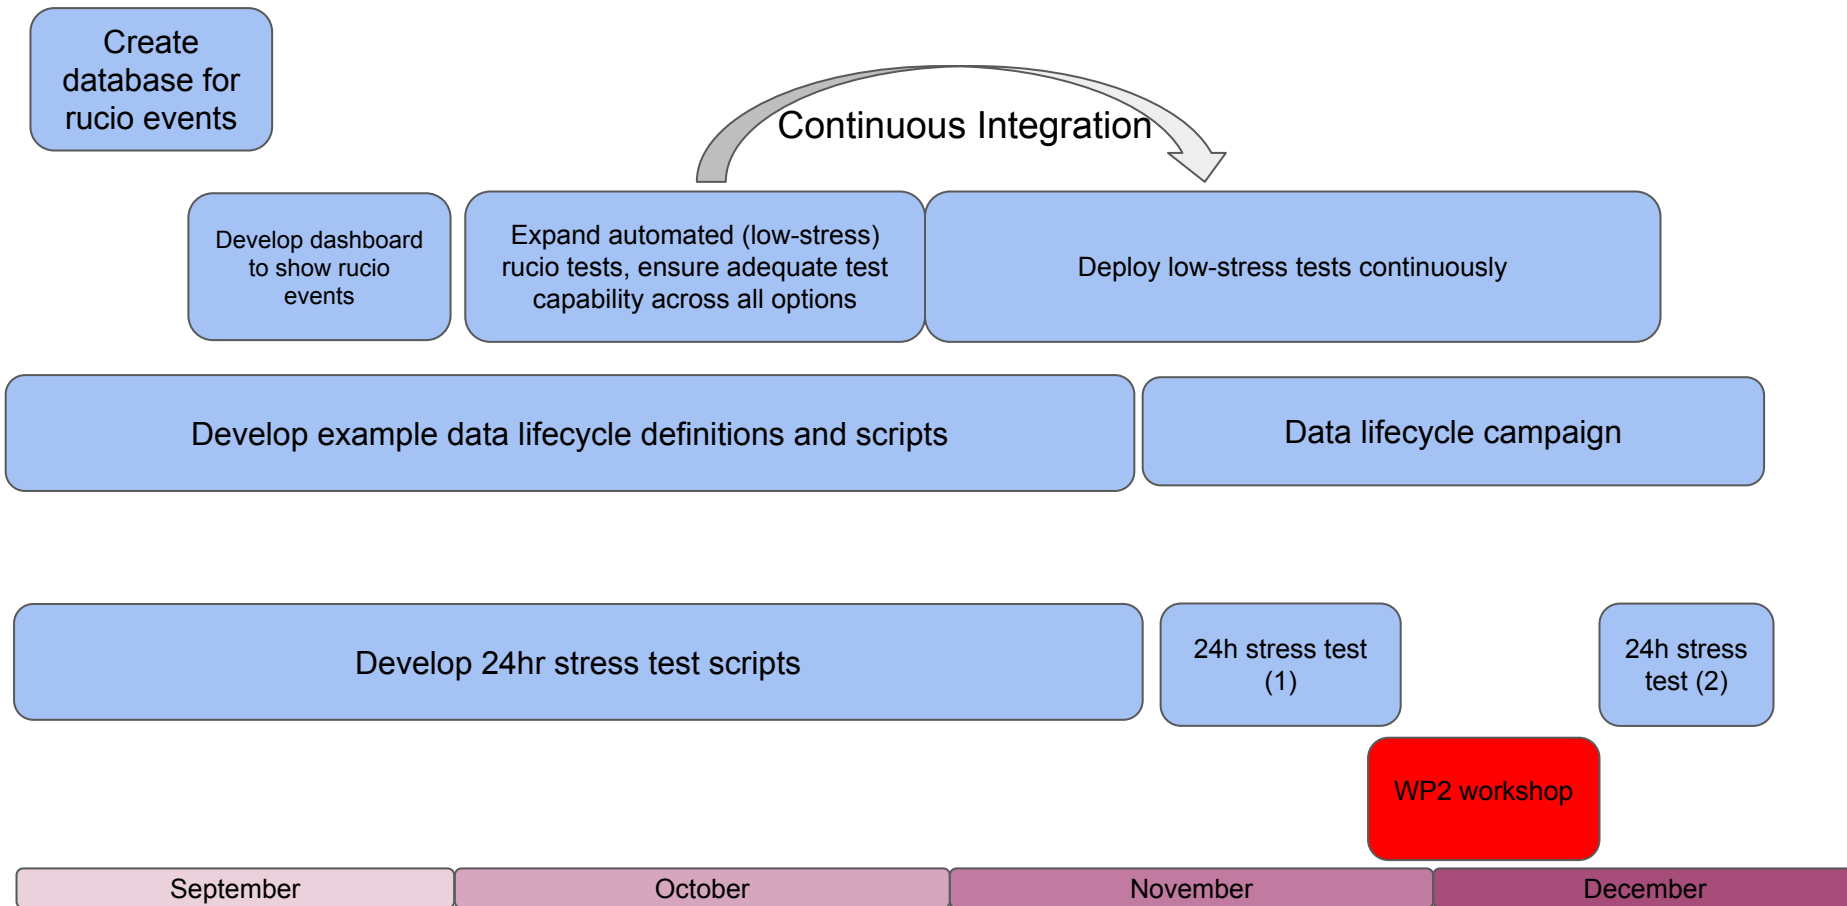
**Goal:** Exercise covering **experiment data workflow** needs on a single day. From data injection, to data replication and data access. Three fold goal: perspective from **scientists**, perspective from **sites**, and the assessment of the **ESCAPE datalake tools and services** under **pseudo-prod conditions**: RUCIO, FTS, CRIC, IAM, PerfSONAR, monitoring, QoS, clients, etc.

**Work plan:** Sept-Oct preparations and tests of the different components in order to run together on a single day, ‘a la’ **dress rehearsal**, mid-November (first challenge, probably a 2nd go afterwards before Xmas break)



1- datalake	2- QoS	3- integration with compute	4- Networking	5 -AAI
-------------	--------	-----------------------------	---------------	--------

# Testing roadmap overview



# Data Lifecycle assessment: Medium-term campaign pilot datalake performance assessment

**Goal:** Test experiment-driven workflows that require a longer campaign - e.g. 4-6 weeks, able to test data lifecycles

**Work plan:**

1 SKAO  
1M files

1 2  
10 sites

1 2  
5 storage technologies

1 4 2  
3 protocols  
http/xroot/gsiftp

2 1  
3 QoS

5 3  
1 AAI  
tokens+x509

Experiment Contacts

5 experiments  
SKA, LOFAR, LHC, LSST, MAGIC

1 2 Experiment Contacts

Data Injection test

2 1

Data lifecycle test

1 3

Data corruption test

4 2 SKAO 1 5

Monitoring test

In-flight transfers, network, storage, files, throughput,...

3 5 SKAO Experiment Contacts

Data access  
(remote I/O, download and through caching layer)

- 1- datalake
- 2- QoS
- 3- integration with compute
- 4- Networking
- 5 -AAI

# Task 2.1: DataLake Objectives and Needs

- 1M files - to demonstrate stable and sizeable data movement
  - Automate data injection from Rucio Client and/or user GET/PUT/POST
  - DataLake accessible by different clients (e.g. simple HammerCloud job)
  - Test data injection and lifecycle using mock data profiting from already existing tools
  - Weekly tests of injection and deletion towards November performance assessment
  - Data corruption starting with “done-manually”
- 10 Sites, 5 Storages Technologies, 3 Protocols
  - Stable infrastructure
  - Transfer matrix green
  - Automated (FTS) tests
  - Automated gfal tests
  - Noise (Rucio) production from defined RSEs

# Task 2.1: DataLake Objectives and Needs

- Monitoring
  - Test in-flight transfers, storages, etc.
  - Rucio events traced e.g. by scope
  - Site/RSE and transfers/configurations carefully monitored
- 3 QoS
  - Starting simply with A, B, and C QoS
  - CRIC as reference point
  - Changing QoS within a site and across sites (on demand, by policy)
- Preparations for WP2-M2.4 (Mar 2021)
  - HammerCloud ready to run realistic research infrastructure workloads
  - Real data distribution and analysis for non-HEP RI (LOFAR, CTA, LSST, MAGIC)
  - Ability to plug heterogeneous clouds (commercial)



# Task 2.2: QoS Objectives

- Demonstrate compute-driven QoS staging:
  - Data is made available for a limited period at resources “close to” some computing resources with the appropriate QoS for that computational work-flow.
- Demonstrate cost-performance QoS trade-off:
  - Data QoS changes so that it is stored on cheaper media.
  - Triggers: elapsed time, inactivity, embargo period ends, ...
- Demonstrate VO-specific work-flows:
  - Exercise selected work-flows from QoS document
- Demonstrate data injection with targeted QoS:
  - Data is written into the data-lake with a desired QoS.
- Demonstrate computational match-making:
  - Select the best computational resource, based on data locality (excluding data locality within any caching and latency-hiding layer) and desired QoS.

## Task 2.3: Compute integration objectives

- Interactive access to files in data lake (e.g. using the Jupyter notebook)
  - Interactive processing by users of experimental data.
  - This also includes transparency of AAI (see slides on that)
- Batch access to files in data lake
  - Simple pipeline(s) to process large amount of data, with fixed parameters
  - Reingesting output data
- Access to data on remote data lake locations through caching layer vs getting data locally for processing
- Data corruption, reporting of corrupted data or fixing it is very relevant, but may be done later on in the project

## Task 2.4: Developing a testing suite

We have already good access to dashboards displaying results from network health monitoring (perfSONAR) and the FTS transfers

- To assess the datalake we need to get rucio events into a dashboard too
  - Urgent work, to be completed by end September, starting immediately
- Develop and run rucio-level tests
  - Use rucio rules, run automatically
  - Test a range of different parameters - file size, file rate, transfer protocol, error recovery, QoS
  - Framework will be the same as for the more intensive tests
  - Build confidence in these tests during Sep-October, before first intensive campaign

# Task 2.5: AAI testbed objectives

- **User enrollment flow in place**
  - Users can enroll in the ESCAPE VO using their institutional credentials, apply for group membership and link X.509 certificates to their account
  - Group administrators for each experiment can approve group membership requests
- **X.509/VOMS AuthN/Z in place**
  - Users can get a VOMS credential reflecting their ESCAPE VO membership and can access data in the data lake using their VOMS credential
  - Users can only access data belonging to their experiment from CLI and Web-based applications
    - Group-based VOMS authorization is in place and correctly configured at datalake sites
- **Token-based AuthN/AuthZ in place**
  - Users can authenticate using their institutional credential and get an access token reflecting their ESCAPE VO membership and can access data in the data lake using their access token
  - Users can only access data belonging to their experiment from CLI and Web-based applications
    - Via scope or group-based authorization

# Status

- We are making quite some progress
- Data ingest experiments are being executed right now “we” (the use cases) are getting our hands dirty!
- Working on deploying the GSOC Jupyter notebook for Rucio on Jupyterhub (SKAO), also a version deployed on SURF infra.
- Defining QOS levels for testing rule-based data movement

# What can you do?

- Rucio - ESAP integration through the Jupyter lab notebook -> Interactive analysis demo
- Batch demo -> We have LOFAR but other use cases very welcome to join in
  - Both for Rucio data access as XCache caching use cases welcome to join
- Especially use cases not really connected to WP2 may be interested because they can provide valuable input from the perspective of first users who are not necessarily interested in implementation details.