

SKA Report on ESCAPE Data Lake Initial Trials

Overview

In preparation for the initial report on the ESCAPE data lake prototyping, ESCAPE partners have been asked to perform some preliminary tests to evaluate the efficacy of the data lake technology stack for their respective use cases. Here we give a brief overview of the testing conducted by SKAO. So far, this campaign has been conducted at the Rucio (data lake orchestrator) level, using randomly generated files, rather than real or simulated data sets.

Environment

SKA testing has been performed using the *rucio-analysis* package, available on GitHub (<https://github.com/ESCAPE-WP2/rucio-analysis>). This offers a Python wrapper around the Rucio CLI, enabling tests to be written in Python and test parameters to be varied in associated configuration files. The tests are run in a Docker containerised environment, which requires an X509 certificate to be mounted in order to make calls to the Rucio REST server. Whilst test development can be performed on any machine with Docker, tests are executed on a dedicated cloud machine hosted by STFC Cloud. Full setup instructions are provided on the GitHub page.

Functional Testing

A functional test (*TestReplication*) of the full grid of sites runs every hour (minute 0) in the Rucio scope SKA_SKAO_BARNESLEY-testing. For each Rucio storage element (RSE), a single 100 kB file is uploaded, and then a replication rule for all other sites is added. This allows identification of sites for which file transfers are currently broken.

Load Testing

In addition to functional testing, a load test has been created (*TestReplicationBulk*), which also runs hourly (minute 10), to test higher velocity data upload and replication. Currently this uploads 1600 15 kB files to the DESY-DCACHE RSE, replicating to SARA-DCACHE and EULAKE-1, in the Rucio scope SKA_SKAO_COLLINSON. The test run time is approx 15 min and increases with the number of files. Thus while it would be possible to increase the scale of this test by a small factor, a larger scale increase could risk the test running for longer than the job interval (i.e. an hourly job taking more than an hour to run). To mitigate this would require

the test to be run less frequently, or a compute instance with more cores being provisioned for the task.

Results

Results can be viewed in the ESCAPE Rucio Events [Grafana dashboard](#). Transfer success ratios are at 100% for the load test (which tests only 3 sites) and typically as follows for the functional test:

Transfer Matrix: transfer-done/transfer-submitted

src	dst	DESY-DCACHE	SARA-DCACHE	PIC-DCACHE	EULAKE-1	LAPP-DCACHE	IN2P3-CC-DCACHE	CNAF-STORM	ALPAMED-DPM	GS1-ROOT	INFN-NA-DPM	PIC-INJECT
DESY-DCACHE		NO DATA	100%	42%	44%	0%	100%	0%	7.5%	0%	91%	60%
SARA-DCACHE		100%	NO DATA	104%	44%	0%	100%	0%	8.9%	0%	79%	100%
PIC-DCACHE		100%	100%	NO DATA	30%	0%	100%	0%	7.8%	0%	84%	102%
EULAKE-1		100%	NO DATA	20%	NO DATA	0%	NO DATA	0%	4.0%	35%	9.1%	NO DATA
LAPP-DCACHE		NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
IN2P3-CC-DCACHE		100%	100%	100%	44%	0%	NO DATA	0%	7.7%	0%	79%	106%
CNAF-STORM		NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA	NO DATA
ALPAMED-DPM		100%	100%	100%	50%	0%	100%	0%	NO DATA	30%	79%	100%
GS1-ROOT		100%	100%	100%	14%	0%	100%	0%	7.4%	NO DATA	100%	100%
INFN-NA-DPM		NO DATA	NO DATA	167%	0%	NO DATA	NO DATA	NO DATA	7.7%	60%	NO DATA	200%
PIC-INJECT		NO DATA	NO DATA	100%	53%	0%	100%	0%	6.0%	0%	NO DATA	NO DATA

Note: some sites show values with greater than 100% success ratios - this is due to requests submitted **outside** of the graphing time interval completing **within** the time interval. This issue is not addressable with the current Grafana/ES combination.

Next steps

The next steps in this data lake testing campaign will be to increase the coverage of functionality we test, including file download and deletion, as well as creating replication rules based on a target quality-of-service, rather than site.