Minutes for 2021-02-03 T2.2 meeting

The previous meeting minutes.

Present

Aleem Sarwar
Berkay Turk
Frederic Gillardo
Federica Agostini
Lucia Morganti
Marek Szuba
Martin Barisits
Pandey
Paul Millar
Rohini Joshi

Apologies

Marcelo Soares (overlap meeting)

Agenda

- ESCAPE wiki pages (Aleem)
- QoS options in the ESCAPE dashboard (Aleem)
- Updates on Rucio developments (Paul)

News

Round table

Streams info can be found at:
https://indico.in2p3.fr/event/19937/contributions/75972/attachments/56927/75787/ESCAPE-T2.2-update.pdf

Stream A: prototyping and demonstration

ESCAPE Wiki: resources page

Aleem has updated the wiki page containing information about the currently available storage options.

It wasn't immediately clear whether all the different storage options are available as RSEs within the ESCAPE testbed.

AP/ Aleem to check whether all storage options listed are available.

Rohini asked whether it makes sense to duplicate information, since some of this information is already available elsewhere; e.g., in CRIC.

After some discussion, we decided to try and push as much metadata into CRIC as possible, including our custom metadata, if that is supported. If this is possible then we would retain the wiki page as the authoritative source of information, but the wiki page would defer the reader to CRIC. This would allow us to annotate the available options and/or include additional (not yet in production) QoS options.

AP/ Aleem to investigate whether we can store custom metadata in CRIC, and populate with our metadata

Possible dashboard improvements

Various possible improvements were discussed.

Active QoS testing

The idea is to provide a high-level verification that QoS is "working". One option is to upload some data into a specific RSE and add a rule that the data should have a specific QoS. The test would pass if the rule is satisfied and the data is stored with the correct QoS. On a practical note, the testing could be done within a specific scope. This would isolate these tests from any other activity.

The Rucio analysis framework (RAF) is already available in ESCAPE github. This already contains examples for similar tests (e.g. uploading, some QoS operations). Please send feedback on the RAF to Rohini and Rob.

The results of these tests could then be fed into the dashboard.

AP/ Rohini to investigate setting up QoS testing.

Usage monitoring

Frederic has demonstrated a jupyter notebook that provided information about QoS usage. We would like to integrate some of these ideas into our regular monitoring.

The rucio admin commands allows querying an RSE to learn which QoS it has. We could then aggregate the total and used capacity of all RSE with the same QoS. This could be plotted as a time-series, allowing the total usage in the ESCAPE testbed to be broken down by QoS class.

For a specific QoS class, the usage may be broken down by VO (or perhaps by scope). This could also be plotted over time, allowing someone to see which VOs are consuming storage capacity with a specific QoS class.

We felt that it wouldn't be useful to discuss this further without the monitoring experts, such as Rizart and Riccardo.

AP/ Aleem to start a discussion on Rocket chat.

Stream B: engagement with experiments

ESCAPE Wiki: the QoS Architecture page

Aleem has created a new wiki page that contains information from the various QoS surveys conducted so far. This page is more than just a series of links to the document; he has copied tables from each ESFRI/experiments documents into the wiki page.

We agreed that the new page is very good, but would benefit from a link to the WLCG DOMA-QoS white paper.

AP/ Aleem to add a link to the WLCG DOMA-QoS white paper.

We discussed the idea of updating the wiki page to include the idea that we include the list of file types that FAIR/CMB provided. CTA has similar information (called Data Levels) that could be included as file types. SKA is (currently) missing this information.

AP/ Aleem to add tables in the new wiki page that includes the different types of file.

Stream C: software developments

QoS Development

Martin reported that one of the ESCAPE fellows (Rizart) will now include Rucio development in his work load. He will likely spend about two weeks taking over responsibility for running the ESCAPE Ruicio instance. After that, he will enable OpenID-Connect (OIDC) support in Rucio. This will be a “dual-stack” solution, where Rucio supports both X.509- and OIDC based authentication. Therefore, the testbed will likely support OIDC by 2021-03-01.

Once OIDC support is enabled, Rizart’s development effort may include implementing missing QoS features. However, his time may be spent on other activities, such as operational issues with Rucio and fixing OIDC problems discovered by ESCAPE usage.

Availability of rucio metadata in ESCAPE testbed

Frederic asked what is the support for metadata within Rucio?

Martin answered that Rucio has a plugin architecture to support metadata queries. One plugin stores metadata using an Oracle database feature that is only available in relatively recent versions of Oracle.

AP/ Martin to check availability of metadata support in the ESCAPE testbed.

Rucio support of metadata indexing

Rohini asked about what support exists in rucio to support indexing.

Martin described how the generic metadata plugin writes the JSON data into Oracle so that (in effect) all fields are indexed. This means that setting/retrieving metadata for a DID will be very fast, but querying for all DIDs that match some metadata predicate might be slow, depending on what metadata exists.

AP/ Martin to check what is supported on the ESCAPE testbed.

Uploading data to RSE expression

Paul asked if the proposed new feature where data may be uploaded targeting an RSE expression (instead of an explicit RSE) has been implemented.

Martin replied: Yes! This is now part of the v1.23.12 release. This new feature was implemented in the client, so it requires an up-to-date client before the feature is available.

The ESCAPE servers are running instance v1.24.2 and the ESCAPE rucio client docker image should be sufficiently up-to-date.

DTNM

The next meeting will be in two weeks at the usual time (2021-02-17T14:00+01:00).