

ESCAPE WP2 Meeting - 22-04-2020

Participants

| | | |
|-------------------|--------------------|--------------------|
| Aleem Sarwar | Ghita Rahal | Paul-Niklas Kramp |
| Alex Clarke | Gonzalo Merino | Raymond Oonk |
| Andrea Ceccanti | James Collinson | Riccardo Di Maria |
| Aris Fkiaras | Killian Schwartz | Rizart Dona |
| Bastian Gounon | Maisam M. Dadkan | Rob Barnsley |
| Berkay Turk | Marcelo Soares | Rohini Joshi |
| Bruzzese Agustin | Marek Szuba | Rosie Bolton |
| Diego Ciangottini | Martin Barisits | Stephane Jezequelk |
| Enrico Vianello | Myroslav Kavatsyuk | Xavier Espinal |
| Fabio Hernandez | Paul Millar | Yan Grange |
| Federic Gillardo | Paul Musset | Zheng Meyer-Zhao |

News

Xavi: ESCAPE Project office is preparing a document about the potential participation of ESCAPE into INFRAEOSC03

Raymond Oonk: Will ESCAPE project have hardware resources to offer?

Xavi: At the moment is not clear, the document with the potential participation of ESCAPE into INFRAEOSC03 document is being drafted, the aim is to put the focus on science driven requirements and needs. The ESCAPE data lake we are developing can be of interest and use.

Xavi: I was asked by the Grid Deployment Board of WLCG to present the ESCAPE project on the 6th of May.

Gonzalo: synergies between the ESCAPE activities and the computing activities towards the HL-LHC in WLCG/DOMA.

Xavier : the overlap is huge and synergies very clear (datalake concepts, RUCIO, FTS). Also WLCG is having soon a review for the HL-LHC computing and preparing documents regarding the work done towards the HL-LHC computing challenge and the forward look.

Rosie: Can the report be shared with the ESCAPE community?

Xavi: Yes, once I get the green light the documents are suitable for public eyes I will share them.

Early implementation and first tests of storage QoS

Aleem, Aris, QoS presentation

https://indico.in2p3.fr/event/21221/contributions/81961/attachments/58367/78339/Test_QoS_endpoints_in_ESCAPE_datalake.pdf

Aleem:

- Overview of Rucio functionality
- Why QoS is important
- The QoS model for ESCAPE Rucio instance
- Demo - Add RSE, Add xroot protocol, Add QoS attribute for newly created RSE, List RSEs filtered by QoS attributes, Test data uploaded to an RSE with a QoS attribute, Create replication rule with destination RSE expression in the form of QoS attribute
- Rucio development identified to make Rucio natively aware of QoS to avoid for example the need for TPC when transitioning QoS for a dataset within the same site

Questions/comments during presentation:

Raymond - systems exist that support multiple storage types on the same path (state of path changes but not the path itself).

Aris - currently Rucio doesn't support this, as Rucio is not QoS aware and it only knows about paths

Paul - yes, we're aware of these issues but would need development effort. Currently working around these limitations

Xavi - we are close to having a prototype which demos what we want to achieve here. Want to be able to inject data with a chosen QoS, which will be allocated to sites accordingly, and this QoS can be updated at a later date. No longer an abstract concept.

Aris - Suggest to have QoS labels that are not tied to a specific technology, and include a default QoS label for sites/data that does not demand a specific QoS

Fabio - How will QoS management differ for different storage systems and will some have native advantages over others?

Paul - dCache provides a REST API for querying QoS of a file and driving QoS transitions. Only the SRM protocol allows the client to (optionally) specify the desired QoS on upload. A file uploaded with other protocols will inherit some default QoS based on the directory into which the file was uploaded (which may be modified later). We would like to fix this, so that files may be uploaded, using a non-SRM protocol, so that they are stored directly with the desired QoS.

During the EU-funded INDIGO-DataCloud project, the CDMI protocol was adopted and the reference implementation updated to support QoS queries and triggering QoS transitions. This reference CDMI implementation used plugins when communicating with storage endpoints. Plugins exist for StoRM and dCache (amongst others), During the EU-funded XDC project EOS implemented the QoS subset of dCache's REST API. This allows the dCache plugin for the

CDMI server to also control EOS. So, in ESCAPE we have the option of targeting CDMI (with the disadvantage of requiring sites to deploy an extra service), or the storage endpoint directly. Martin - There should be native QoS support from Rucio for different protocols including webdav, if not now in the future

Shadow Round Table

Aris: PIC noticed that some of the TPC transfers were failing. Found an issue with DCache (failing to create parent directories). Unclear the status of this - recommends updating DCache.

Paul: Aris is right, there is an issue with DCache. Issue seems to be evolving, not fully understood. XRootD server creates parent directories irrespective of TPC, but shouldn't by default for direct transfers (not TPC). Not established what the correct behaviour should be.

Aris: Did not see this issue previously, must have been introduced recently.

Paul: Haven't seen anything that should trigger this regression. There is HTTP/TPC is available as an alternative.

Xavi: Any other site updates?

Rohini: We'll be looking to upload some SKA data to ESCAPE data lake at some point later this week. Site question as follow up to Aleem presentation: Is there a way to have an RSE expression with an upload, to specify a RSE?

Martin/Aris: Not really - but you upload to a site, and then specify QOS.

Rohini: One more thing; getting people on our side (SKA) more familiar with Rucio. Using elements from Agustin's python script and developing a Jupyter notebook.

Xavi: Action point (8/04 TBD Ghita and Paul Musset: Update on XCache project@IN2P3 at the WP2 meeting on the 13th of May (tbc)) still looking good.

Rizart: Still looking for input on monitoring requirements.

Xavi - Next meeting 6th May