

DOMA-QoS 2020-11-25

Present

Paul Millar
Martin Barisits
Chih-Hao Huang
Mario Lassnig
Marcelo Soares
Peter Love
Oliver Keeble
Duncan Rand

Agenda

Progress on demonstrators
Fallout from WLCG meeting

Demonstrators

After a discussion, we agreed to write a template for describing demonstrators in general, which may be used as a starting point for all QoS demonstrators. The idea is to provide some documentation about the components of the demonstrator, what makes this demonstrator special, which work-flows will be executed, how the demonstrator will be evaluated, etc.

From these templates, we will create two specific instances: one for the HPC / Opportunistic-storage demonstrator (Doug's demo) and one for the Edinburgh/Opportunistic-storage demonstrator (Rob's demo), with the blank spaces filled with our best knowledge (or place-holder text). We would invite Doug and Rob+ATLAS to correct any mistakes and fill in any parts that are missing.

Mario mentioned how there is a JIRA ticket with plans for the Edinburgh: [ADCINFR-184](#)

This ticket has been appropriated by LRZ, which is providing some similar storage: out-of-warranty hardware. The capacity is currently some 200 TB, with up to 600 TB by the end of the year. This is represented in Rucio as a new RSE: LRZ-LMU_VOLATILE.

There was a discussion on whether this storage is raw (plain JBOD) or deployed with some redundancy (e.g., RAID). After consulting with the ticket, the plan appears to be to avoid redundancy and use the available storage hardware to give the maximum available capacity.

Mario also observed that the LRZ-LMU_VOLATILE RSE currently has 2 TB storage in use, with 20 TB as the configured limit: [DDM-MON: LRZ-LMU_VOLATILE](#). That 2 TB usage breaks down by scope as:

scope	nr_replicas	sum_bytes
mc15_13TeV	1	94834717
mc16_13TeV	419	2396203598171

AP/ Paul to drop Chris an email to see if we can “appropriate” LRZ-LMU_VOLATILE as a new demonstrator for DOMA-QoS.

We also have our potential contact with Frank Wurtwein, which unfortunately we haven't followed up. In the WLCG workshop, Frank [presented](#) his strategy for using out-of-warranty hardware in a JBOD deployment.

AP/ Paul to drop Frank an email to see if we can add this JBOD as a QoS demonstrator.

WLCG workshop

Oliver provided a brief summary of the WLCG workshop. He described how there were talks about “analysis facilities”, where one distinguishing feature of these facilities is that they are provisioned with special, low-latency storage. These facilities are not really “on the grid”, but there were some discussions on how they might be added, with some protection to ensure that only the desired work-flows are accepted.

There was a talk from our friends at KISTI, describing their novel archive solution. They presented this at our QoS workshop.

One point of note was a discussion over a very specific problem that LHCb has with transferring data from CTA. This discussion evolved into a discussion about how tape might be managed via HTTP. The discussion dwelled briefly on the idea of adopting the dCache REST API as a starting point for a new WLCG-wide tape storage API. Memories of SRM and incompatible implementations stimulated a desire for any such API to be minimal.

Further discussion on this topic has been scheduled for some future (unspecified) DOMA meeting.

Supporting Rucio development

Rucio is developing support for VO policies. This will be ready in February at the earliest. In the meantime, we (the ESCAPE project) would like to support Rucio development by extending the ESCAPE testbed to include CDMI servers and the ability for Rucio to trigger QoS transitions.

Although the development was completed within the XDC project, it is unclear whether this will work out-of-the-box. There may be the odd problem that comes to light once FTS+CDMI+storage is tested. Given the likelihood of dropping CDMI, it is unlikely that fixing these problems will be a priority for the FTS team.

Olivier expressed his confidence that Rizart's skill-set would allow him to fix problems with the FTS-CDMI interface. So this might be a way to get an ESCAPE CDMI demonstrator working.

AP/ Paul will approach ESCAPE to see if this is something we could ask Rizart to invest some time, at least for deploying/updating FTS to support QoS, configuring FTS, and to help in diagnosing problems.

After a little discussion, we agreed that DOMA-QoS is interested to learn more about the ESCAPE testbed's use of CDMI, but would not contribute to this directly.

DTNM

The next meeting will be in two weeks time, at the regular time slot: Wednesday 2020-12-09 at 16:00 CET.