

From: Alastair Dewhurst - UKRI STFC alastair.dewhurst@stfc.ac.uk

Subject: CTA Evaluation

Date: 13 September 2019 at 17:40

To: Steven Murray Steven.Murray@cern.ch, German Cancio Melia German.Cancio.Melia@cern.ch, Julien Leduc Julien.Leduc@cern.ch

Cc: Alison Packer - UKRI STFC alison.packer@stfc.ac.uk



Hi all

As you are aware RAL has been looking at replacement for Castor. For various reasons we had been looking at a commercial solutions. A tender was launched in August that recently failed without attracting any bids. Rather than re-launch the tender again with more money we decided to re-evaluate CTA first.

We (Alison and I) have been speaking to Alberto Pace and will be in CERN on the 3rd - 4th October to discuss some of the strategic level questions with him. We would also like to have a discussion with you guys on the more technical side of things. Are you available either of those days? I am out at an ATLAS C&S week, so I could come as see you on Wednesday if necessary.

We would like to understand:

1) What would it take to setup and run a RAL sized CTA instance?

We are very keen to run a service that is as close to what CERN are doing as possible. We really want to minimise operational risk/complexity, so we would run a small EOS instance etc if it was necessary.

We currently run 2 production instances, WLCGTape and Facilities. The WLCGTape unsurprisingly hold the data for the WLCG experiments as well as a handful of other HEP experiments. It has about 60PB of data stored. It has a ~2PB shared buffer and workflows that are pretty similar to what CERN sees. The Facilities instance has a similar amount of data stored and is primarily used by local (i.e. RAL) experiments. There are several services built on top of Facilities Castor that provide individual users with the ability to access their data.

Do you have plans detailing how you setup new clusters?

2) How would we migrate data from Castor to CTA?

I am sure you have plan for how to do this!

3) What are the long term development goals

We have heard that you are working on using postgresSQL databases rather than Oracle. Anything else exciting come up?

Thanks

Alastair

P.S. I can see Julien is giving a talk on 27th September on CTA (<https://indico.cern.ch/event/843592/>), I assume these have a webcast?