

From: Michael Davis michael.davis@cern.ch
Subject: Re: CTA deployment meeting, tomorrow 11/7 @ 14:00 in 31/S-023
Date: 11 July 2019 at 09:21
To: German Cancio Melia German.Cancio.Melia@cern.ch, castor-tape-dev (castor tape developments) castor-tape-dev@cern.ch



Update on action points for migration test:

1. Compile EOS version with required changes: gRPC API, new checksum protobuf format, prepare request ids/ABORT_PREPARE. **DONE**
 - **Tested Prepare Cancel API. Not yet implemented in CTA, but a log message is generated by the CTA Frontend.**
 - **There is an outstanding XRootD bug when sending Prepare with no parameters, CTA handles this situation gracefully.**
2. Merge CTA schema changes into master. **DONE**
3. Review final DB schema for migration.
 - **I propose we review this at our next deployment meeting on 18/07/2019, in order to confirm the final schema with Giuseppe when he returns (22/07/2019).**
 - **We will use the currently-deployed schema for the tests w/c 22/07/2019.**
4. Create DB migration tools for ATLAS instance: alter schema and convert checksums and uid/gid to new format. **DONE**
 - **Julien and I updated the schema on the ATLAS instance and deployed a CTA version that uses the updated schema (10/07/2019). We tested that existing files can be retrieved and that new files are archived and retrieved correctly. We will continue to monitor during ATLAS data taking today.**
5. Update EOS namespace injection tools to use new gRPC API.
 - **EOS with gRPC API is deployed on ATLAS instance. We need another instance for the migration test.**
 - **Giuseppe has written SQL scripts which create the intermediate tables required for injection into EOS namespace.**
 - **I am currently updating migration tools for directory and file injection.**
6. Small-scale metadata migration to validate all tools and workflow for the migration, including handling failure modes.
 - **requires migration tools above, deadline 19/07/2019**
7. **Milestone:** CASTOR DB to be moved to new hardware. Propose to move CTA ATLAS DB during the same maintenance window. Date to be set by DB team, I believe it is going to be around 15 July, Giuseppe is coordinating with DB team.
 - **Upgrade of CTA DB hardware DONE (Thanks Julien!)**
 - **Migration of CASTOR to new hardware postponed until 22/07/2019 when Giuseppe returns.**
8. **Milestone:** Week 22-26 July: Full-scale ATLAS migration test (metadata only, no tapes). This is a functional and performance test. It will allow us to accurately estimate the time needed to do the real migration and to consider if we need to make any further optimisations.
 - **We are on track to be able to perform this test during w/c 22 July.**
 - **Note that this test is metadata-only (no tapes, no recalls, files in CASTOR will not be marked as exported).**
 - **Tests with tapes/recalls will be carried out in August in order to be ready for migration in September.**

QUESTIONS FOR DISCUSSION

1. How will we import Storage Classes? Do we import CASTOR fileclasses "as-is", or will there be some mapping of CASTOR fileclasses to CTA storage classes (renaming, consolidating several storage classes into one, etc.)?
2. How will we set up Archive Routes in CTA?
3. How should CASTOR symbolic links be imported?

QUESTIONS FOR GIUSEPPE

1. Some directories in CASTOR set SGID, SUID, SVTX (sticky) bits. Does CASTOR honour these bits and should they be imported into CTA+EOS? (This question does not affect ATLAS, just some of the smaller experiments and a few users).
2. EOS does not have ATIME. Confirm how CASTOR file/directory timestamps should be imported.

From: Michael Davis michael.davis@cern.ch

Subject: Re: CTA deployment meeting, tomorrow 11/7 @ 14:00 in 31/S-023

Date: 11 July 2019 at 11:26

To: German Cancio Melia German.Cancio.Melia@cern.ch, castor-tape-dev (castor tape developments) castor-tape-dev@cern.ch

MD

Additional question:

4. Which EOS extended attributes should be set when migrating directories? Minimum is to set CTA_StorageClass which is set on a per-directory level. CTA_TapeFsid is always 65535 so should probably be set by an attribute link. Then there are the EOS system attributes such as:

```
("sys.acl","u:100:rwX")
("sys.forced.checksum","adler")
("sys.forced.space","default")
("sys.forced.nstripes","1")
("sys.forced.layout","replica")
```

On 11/07/2019 09:20, Michael Davis wrote:

QUESTIONS FOR DISCUSSION

1. How will we import Storage Classes? Do we import CASTOR fileclasses "as-is", or will there be some mapping of CASTOR fileclasses to CTA storage classes (renaming, consolidating several storage classes into one, etc.)?
2. How will we set up Archive Routes in CTA?
3. How should CASTOR symbolic links be imported?

