



# ESCAPE

European Science Cluster of Astronomy &  
Particle physics ESFRI research Infrastructures

## Progress report for QoS in DAC21

Paul Millar <[paul.millar@desy.de](mailto:paul.millar@desy.de)>



# Milestone 1: The use-cases have been identified

- We currently have information about DAC21 from SKA, CTA, FAIR and LOFAR.
  - FAIR DAC21 use-cases
  - CTA DAC21 use-cases
  - LOFAR DAC21 use-cases
- Documents cover all DAC21 activity, not just QoS-specific.
- Information is rather scattered at the moment: should we consolidate this?
- We're currently missing any indication of how much capacity is needed.
- Interesting additional storage have been identified:
  - SKA – object storage / CEPH
  - CTA – cloud storage / S3



# Milestone 2: resources needed are deployed

- Mostly focused on making tape available
  - PIC have already volunteered, but not yet made available
  - Negotiated tape access at DESY, with some caveats.
- SKA (Rohini?) and CTA (Frederic?) are organising the required capacity for Ceph and S3 storage.



# Organised tape access at DESY

- ESCAPE has no hardware budget – so we need to be “inventive”
- Negotiated with DESY operations teams that ESCAPE has access to “spare” tapes
  - DESY is willing to run a reduced spare capacity until DAC21 is over.
  - ~100/200 TiB capacity available.
- DESY needs those tapes back – they’re only “on loan”.
  - We can have them for DAC21, but the stored data will be lost after that.
- So ...
  - Don’t store data you care about!
  - We should probably think what is the best procedure for removing data/replicas.



# Tape commissioning process @ DESY

- Enable tape support in underlying storage & test it works! (2--3 days)
- Identify and create directory that we wish to use for Rucio (1 day)
- Update dCache configuration so files in target directory end up on tape (1 day)
- Update CRIC configuration, adding new tape-based RSE into Rucio (1 day)
- Testing that Rucio can write into tape (1-2 days)

Provided there are no unexpected show-stoppers, the process should be complete by end of August.

**We're looking for volunteers who are willing to help out with testing.**





# Tape commissioning process @ PIC

- Probably follow a similar process as with DESY.
- Expected lead-time of ~2 weeks before storage is available for testing.



# Open questions

- How realistic should the tape activity be?
- Tape drives:
  - Normally, tape drives are shared among users.
    - DESY has some “big hitters” with sporadic access patterns.
    - If they coincide with DAC21 then the performance will suffer.
  - We may be able to reserve some tape-drives for the duration of DAC21.
    - +ve side: tape access would be “faster”
    - -ve side: might not best reflect real-life
- Disk cache: how much do we need?
  - how much data is going to be written and recalled?
  - Should we “wipe” the disk cache, to force tape staging?



# Other resources ...

- CEPH storage: would likely need to coordinate the QoS labels
- Is this SKA-only storage? How do we do that?
- Similarly for the cloud (S3) storage for CTA: which QoS label? Is it CTA-only?
- Is there interest in trying out erasure-coding (EC) storage as a QoS class?
  - Option offered by CERN / EOS.





# Active QoS testing

- Ticket <https://jira.skatelescope.org/browse/EDLK-116>
- The plan is to provide some high-level verification that QoS is “working”.
- Much of the initial work has been done as a lead-up to the FDR
- There is some more work still needed.



# QoS monitoring

- Ticket <https://jira.skatelescope.org/browse/EDLK-115>
- Excellent work from Alba.
- Work is now mostly complete, but some (minor) integration into the monitoring framework is still needed.



# Next steps

- Finish commissioning available tape storage
- Work with SKA (Rohini?) and CTA (Frederic?) to ensure ceph / cloud storage has appropriate QoS classes.
- Work with ESFRI communities to start writing and testing software to drive the DAC21 use-cases: injection, replication, processing. (to achieve Milestone 3)



- Milestones (Further details in [this Google docs](#)):
  - M1. **The use-cases have been identified** (due end of May)

All ESFRI communities have a detailed plan on what operations they would like to demonstrate. Any missing features or resources are identified.
  - M2. **The resources needed to support QoS use-cases have been identified and deployed** (due: end of June)

Sites have been identified that will provide access to this hardware, and has been deployed.
  - M3. **The code needed to drive the use-cases is feature-complete** (due: end of July)

This milestone is about developing the code to drive DAC21 QoS and caching use-cases and work-flows. The code is not expected to be feature-complete but not fully tested.
  - M4. **All the QoS use-cases have been exercised at small-scale** (due: end of August)

This milestone is about catching the more obvious problems.
  - M5. **The QoS use-cases have been demonstrated at the desired scale** (due: end of September)

Verify the work-flows work at scale; may require coordinated access.
  - M6. **DAC21 has completed successfully.** (2021-11-xx)

