



ALICE data management considerations

2016-10-08

Maarten Litmaath

v1.0

Run 2



- Only adiabatic changes
- T2 sites should continue to provide fairly reliable storage
 - Pure XRootD is not complex
 - EOS slowly on the rise (9 sites + CERN)
- All storage has to be accessible through the Xrootd protocol
 - And support ALICE envelopes
- Further analysis efficiency improvements being looked into
 - Analysis train fraction to be increased further
 - AOD vs. ESD fraction ditto
 - Nano AOD formats

Accessing data



- Jobs get a sorted list of SEs per file, for reading or writing
 - Based on network topology and SE availability
- Remote access can be tolerated at a level of ~15%
 - Higher use would impact networks as well as the disk servers
- T2 storage as cache, i.e. an SE that may lose data
 - Jobs already fail over to other replicas as needed
 - A “self-healing” cache in a multi-site storage would be opaque and therefore OK
- Regional federations
 - Run 2: OK if opaque (cf. NDGF-T1)
 - Run 3: explicit building blocks of the grid

Clouds & monitoring



- Cloud storage
 - Could be OK with Xrootd interface (cf. DPM in T-Systems)
 - To be able to use pure S3 etc. would require major development
 - Network topology and bandwidth considerations
 - A VOBOX should be located close to the storage
 - The easiest is to use cloud resources only for MC simulation!
- Monitoring
 - ALICE have sophisticated storage monitoring down to individual disk servers
 - I/O, occupancy, system parameters, system health
 - Comparisons with the file catalog allow dark data to be located
 - Data popularity tracking allows unused data to be identified
 - Candidates for removal or replica reduction

Run 3



- Analysis in dedicated Analysis Facilities
- T2 sites are organized in regional “clouds” that manage themselves
 - E.g. running a distributed, self-healing storage infrastructure
 - With fast networks interconnecting the constituent sites
- New technologies for scalable catalogs being studied
 - E.g. Cassandra + CVMFS, in collaboration with Openlab
 - EOS Diamond?