

DDM ops report

ATLAS Software & Computing Workshop

Cédric Serfon

Ludwig-Maximilians Universität, München

April 4th 2011

DATADISK consolidation

- All datasets on T2_MCDISK migrated in January (subscriptions). Sites removed from TiersOfATLAS beginning of February. Squads were asked to inform sites that Space Token can be decommissioned and all files under atlasdisk could be removed.
- For T1_MCDISK, migration done via subscription except for RAL-LCG2 were files migrated manually from MC to DATADISK + LFC updated (thanks to big help of UK squad).
- During the migration a lot of incomplete datasets identified with files that cannot be found anywhere (probably due to old bugs/task migration).
- Only 3 sites still to be migrated : BNL-OSG2, IN2P3-CC, TRIUMF-LCG2. Only BNL should take more than one month to finish the migration.
- See Ueda's talk for more infos.

Lost/corrupted files

- Squads are now informed daily for suspicious bad files.
 - Information collected from the dashboard, and by developpement version of the pilot.
 - That's important that the squads confirm if these files are really bad or not.
- Number of bad files reported since January 1st : ~120k (about 44k definitively lost)
- Some recurrent problems spotted :
 - Nordugrid : Corrupted files due to bad cluster.
 - Taiwan : Seems to be affected by the "RAL syndrom" (bad batch of pools).

AKTR/prodsys issues

- A big part of the Savannah bugs reported are linked to prodsys/AKTR.
 - Non frozen datasets.
 - Aborted datasets.
 - Empty containers.
 - Group datasets.
- Cross check would be needed between prodsys and AKTR info to have the complete picture. Some ideas to improve the issues. To be discussed this week.

Deletion issues

- Actually many deletion issues related to LFC :
 - Deletion suffered problems in some US T2s because of slow LFC deletion (no bulk deletion methods used because of old LFC version). Now every sites have installed the version that support bulk methods.
 - In the last weeks 1000s of deletion requests sent by Borut and group people → big deletion backlog. In big clouds, despite of bulk deletion, LFC deletion seems to be the bottleneck now. Reason : only one catalog cleaner per LFC vs one storage cleaner per SE endpoint.
 - Recurrent ACLs problem in LFC due to old bugs :
 - Manual fix. But hard to do it since files can be owned by 4-5 different people (pilot factory owners).
 - One central LFC + unique service certificate for Panda and pilot factory would help a lot.
- Some race condition panda/deletion leading to missing files at PRODDISK (won't happen if deletion was instantaneous. Adcos shifter are now taught to ignore them.

User support

- First level of user support goes through DAST that does a good job.
- But many questions wouldn't need to be sent to the DAST since already explained in the Wiki (RT*M !) or already answered before in another thread on the DAST mailing list.
- Maybe some improvements are needed :
 - Improve user documentation in the Wiki.
 - More explicit error messages pointing unexperienced/lazy users to the documentation (e.g. cannot find any sites not blacklisted, see <http://blahblah> for more information).
- DATRI problems represent a significant fraction of the Savannah submitted and are sometimes difficult to solve without Mikhail's help. What's the longer term plan for DATRI ?

Miscellaneous

- Some squads sometimes not very responsive to tickets. Would be good if they can be more (pro)active.
- Some DQ2 commands takes now very long and time-out. Need to identify all the queries with this behaviour and report to DQ2 dev for correction (longer timeout or change at the DB level).
- Since we've started the sonar tests, weekly bursts of file exists due to failed FTS attempts in previous sonar tests campaign. Site Services that submit these tests now have option overwrite (but doesn't seem to be working properly).

Next months plans

- LFC migration (see Graeme's talk in DDM session).
- Groupdisk merging.
- More T2Ds.
- User quota on SCRATCHDISK probably need to be rediscussed : last week once again one example of single user writing 5-6 TB to a site in less than 1 day !
 - Victor is not fast enough to handle this kind of issue.
 - The files from this user will stay at least 2 weeks w/o being deleted. This can block completely the site.