

FTS-ATLAS@CERN issues

Andrea Manzi
on behalf of the FTS team

FTS scheduler slowdown

- Starting from the 7/08, ATLAS started a tape recall campaign from T1s (and new EOS-CTA @CERN)
- The number of requests sent to the FTS ATLAS @CERN was quite high > 1M staging+ transfers
- After a couple of days of normal operation on the 9/08 in the early morning our daily process cleaning the DB from older requests started and it dramatically loaded the FTS DB, so all scheduler queries got slowdown (10% of the previous performances)
- The issue lasted most of the 09/08 and again from the 10/08 early morning till the afternoon
- [GRAFANA link](#)

FTS scheduler slowdown: consequences

- The main consequence of this problem was the IN2P3 poor staging stats, why?
- We asked IN2P3 to recall 700 TB of data, this went pretty smoothly but FTS was not able to schedule transfers between the tape buffer to the disk on time
- When the Scheduler performances went back to normal the files FTS tried to copy were already garbage collected from the Tape buffer (the size there is 110TB)
- Not sure about the exact numbers but I think ~400 TB needs to be restaged because of this issue

- On the 22nd we have also an unusual problem with the staging polling (some requests got lost on a couple of servers after a crash), restarting the daemon solved the issue and we understood how to fix it.

How to improve things?

- We have a new version of FTS (3.9) ready to be installed (already running in the pilot and at RAL) which is improving the DB schema and the scheduling performances under load (also with our DB cleaning process running). We just need to agree on a date for the upgrade as it needs a downtime
- We are working on the usage of table partitions to improve the DB performances. We are also discussing with IT DB to have a dedicated DB node for us which we can better tune. We now run using on a shared DB instance via DBonDemand and we cannot change some of the parameters there.
- Rucio will also throttling the staging request to FTS which are the ones staying longer on our DB and degrading the performances (under implementation)
- Ask Rucio to move some links to other FTS servers, as a temporary measures, as our instance gets the majority of the load.
- Improve the monitoring/alarms on our side to understand when issues like this are happening, we only understood the source of the problem days later.