



WLCG Service Report

Jamie.Shiers@cern.ch

~ ~ ~

WLCG Management Board, 10th February 2009

Overview

- Today's report will be short: no "major" incidents last week
- There was the usual background of problems that were addressed as they arose
- Some "SIRs" still outstanding:
 - Interim report from FZK on Oracle problems affecting FTS/LFC
 - Report on CMS 500 lost files at FZK
 - [Also] high failure rate seen with ATLAS transfers to FZK (20-80%) – problem understood and requires a compact of DB

GGUS Summary

VO concerned	USER	TEAM	ALARM	TOTAL
ALICE	3	0	0	3
ATLAS	25	8	0	33
CMS	5	0	0	5
LHCb	11	2	0	13
Totals	44	10	0	54

- Nothing special

FZK DB SIR

- At GridKa/DE-KIT the FTS/LFC Oracle RAC database backend was down from January 24 to 26 (Sat, approx. 0:00 to Mon, approx. 22:30 CET).
- On Sat, our on-call team immediately received Nagios alerts. From approx. 9:30 on Sat, our DBA worked on the issue and found, that many Oracle backup archive logs had been filling up the disks.
- By trying to add an additional disk, ASM (the Oracle storage manager, i.e. the file system) got blocked. The reason was probably a mistake made by the DBA when preparing the disk to be added.
- Due to the fact, that the LFC data was on the affected RAC system and it was unclear if the last daily backup worked properly, the DBA decided not try simple repair attempts like rebooting nodes etc but to involve Oracle support.
- At approx. 16:30 on Sat. she opened an Oracle Service Request. After info/files exchange with an Oracle supporter (in timezone CET-8h) till Sat late night, another supporter (in our CET zone) came back to us on Mon, approx. 11:00. With his aid, the problem finally was solved.
- Remarks:
 - It is unclear to me, why it took more than a day until we got an Oracle supporter in our timezone. It could be, that the support request was not filled in correct. I wanted to clarify this before sending a SIR since it is not clear if bashing on Oracle is fair in this case.
 - As soon as I get to talk to the DBA I will try to clarify on which side mistakes happend.
 - My personal opinion: even though the disk to be added to the ASM was not prepared correctly, the system should not block but the command issued to add the disk should throw an error message. From a software costing thousands of Euros per licence, I would expect that.

FZK dCache

- FZK will be able to do this not sooner than Tuesday. The operation means taking down the dcache nameservice (pnfs) for all supported vos.
- We have to discuss this with them. The ATLAS table space has grown to disproportionate size because deleted records still linger. The quickest way to compact the db is to dump and restore it which takes an expected 5 hours in total. It may cure the problem. It may not.
- It will not work miracles. The SRM is still under high load and no medicine has been found yet to heal it. We most likely will dump/restore on Tuesday or Wednesday but can tell you definitely no sooner than Monday afternoon.
 - FZK running with 20-80% failure for file transfer / stagein-out. This comes from pnfs overload, due to amount of files and problems with DB load. Would be cured with ~1/2 day operation on pnfs database. FZK would like to wait ~1 month when they will split SRMs per vo. This would mean fairly degraded mode for 1 month. Rod will push for this 'healing' done asap and not in one month.

FZK – CMS lost files

- *"a few days after the problem had been discovered, we were sure that we had identified all affected files (actually we could recover many of them at GridKa)"*
- *"However, I forgot to report this incident later to the Ops meeting, after it was fully understood. Sorry for that."*

--Andreas

More...

That's all for this week's service report

- I have a couple of questions related to the workshop...
 1. Given the recent LHC news does it make sense to have a talk from LHC operations or should we save this for later?
 2. Experiment plans – sites often say that they are not clear in terms – could a site provide a rapporteur to collect and present the information in the right format?
 3. More suggestions – they typically come in the last days (or hours...) which can be difficult to handle...