



WLCG Service Report

Harry.Renshall&Jamie.Shiers@cern.ch

~ ~ ~

WLCG Management Board, 27th January 2009

GGUS Tickets submitted concerning WLCG VOs from 2009-01-18 00:00:00 thru 2009-01-24 23:59:59

Tickets affecting WLCG VOs

	USER	TEAM	ALARM	Total per row
atlas	14	17	3	34
cms	0	0	1	1
lhcb	3	1	0	4
Total per column	17	18	4	39

Summary of the alarm tickets:

CMS to CNAF created 20 Jan, addressed same day and marked as solved 21 Jan

ATLAS to FZK created 21 Jan and marked as solved same day

ATLAS to RAL created 22 Jan marked as verified same day

ATLAS to CERN created 23 Jan marked as verified same day

Experiment Alarm "Tickets"

- CMS to CNAF:
 - Tue 20 Jan D.Bonocorsi reported Phedex exports from CNAF failing since a few hours with Castor timeouts.
 - Additional error was reply he was not allowed to trigger an SMS alarm for INFN T1 (being followed up)
 - Problem addressed same day and found to be due to a single disk server.
- ATLAS to FZK:
 - Test of alarm ticket workflow after new release. Closed within 1 hour.
- ATLAS to RAL:
 - Thur 22 Jan RAL failing to accept data from Tier 0 giving an Oracle error on a bulk insert call. Within 1 hour solved by restarting SRM processes after which FTS reported no further errors.
- ATLAS to CERN:
 - Fri 23 Jan from 15.00 almost all transfers to CERN ATLASMCDISK space token failing with 'possible disk full' errors. This was due to a misconfigured disk server that was then removed. Then on Saturday importing failed again which was reported back as the pool being full. Monitoring showed it full while stager and srm queries did not. Removing the misconfigured disk had also taken out state information, a known castor problem. The machine was back in the pool Monday and the state information resynchronised.

Other Outstanding or new “significant” service incidents (1 of 3)

- ASGC: Jan 24 FTS job submission failures due to constraint of ORACLE maximum tablespace. Had to add 100MB manually. Follow-up is to try adding a new plugin to monitor the size of table space to avoid the same situation in the future.
- FZK: Jan 24 The FTS and LFC services at FZK went down due to a problem with the Oracle backend. The problem was quite complex and Oracle support was involved. Reported as solved late Monday.
- BNL: Jan 23 hit by the FTS delegated proxy corruption bug, a repeated source of annoyance. Back porting of the fix from FTS 2.2 to 2.1 is now in certification and eagerly awaited.

Other Outstanding or new “significant” service incidents (2 of 3)

- PIC: Jan 24 Barcelona was hit by a storm with strong winds and suffered a power cut about midday which turned off air conditioning and closed them down. Resumed Monday morning being fully back around midday. Had some problems bringing back Oracle databases following their unclean shutdown.
- CERN: Jan 22/23 the lcg-cp command (which makes srm calls) started failing when requested to create 2 levels of new directories. Recent SRM server upgrade suspected - to be followed up.
- PIC: Jan 21 high srm load due to srmLs commands thought to be from cms jobs (as previously seen at FZK). No easy/feasible control mechanism.

Other Outstanding or new “significant” service incidents (3 of 3)

- CERN: Jan 20 25K LHCb jobs "stuck" in WMS in waiting status. Further investigations suggested an LB 'bug' that occasionally leaves jobs in limbo state. Plan is to see if latest patch will fix. In the meantime DIRAC will discard such jobs after 24h.
- CERN: Jan 19 ATLAS eLogger backend daemon hung and had to be killed. Follow up will be to trap the condition and write an alarm. Heavily used by many levels of ATLAS and they did not know where to report it. Follow up has been to create a new ATLAS-eLogger-support Remedy work flow with GS group members as the service managers.

Situation at ASGC

- Jason was at CERN last week and we (Jamie) discussed F2F issues around their DB services –3D & CASTOR+SRM and including migration from the OCFS2 filesystem.
- A target date for restoring the “3D” service (ATLAS LFC & conditions, FTS) to production is early February 2009
 - The new hardware should be ready before that allowing for sufficient testing and resynchronization of ATLAS conditions then VO-testing before announcing the service to be open
- A tentative target for a “clean” CASTOR+SRM+DB service is mid-February 2009 – preferably in time for the CASTOR external operations F2F (Feb 18-19 at RAL)
 - It is less clear that this date can be met – need to checkpoint regularly (work will start after Chinese NY on Feb 2 – see notes)
- ASGC will participate in bi-weekly CASTOR external operations calls & 3D con-calls + WLCG daily OPS

Other Service Changes

- GGUS has now released direct routing to sites for all tickets not just team and alarm. Submitter has new field to target site when it is known.
- CERN moved back into production the two LCG backbone routers that were shutdown before Christmas (attended by hardware support engineers).
- CNAF been testing long-term problematic access to the LHCb shared software area in GPFS. A debugging session discovered that the GPFS file system, for intrinsic caching mechanism limitations, runs the LHCb setup script in 10-30 minutes (depending the load of the WN) while a NFS file system (on top of GPFS storage) runs the same in the more reasonable time of 3-4 seconds as at other sites. A migration is planned.
- IN2P3 following the ATLAS 10Million file test have deployed a third load balanced FTS server addressing CNAF, PIC, [RAL](#) and ASGC. Some improvement though not much was initially seen.

Summary

- **Business as usual - the rate of new problems arising and that need follow-up service changes remains high.**
- **The problems remain heterogeneous and without a predictable pattern to indicate where effort should best be invested.**
- **However, the resulting service failures are usually of short time duration.**