



WLCG Service Report

Jamie.Shiers@cern.ch

~ ~ ~

WLCG Management Board, 2nd December 2008

Summary of the week

- **ASGC – continuing problems with CASTOR service. Current issues are clearly with Oracle service / configuration. (From last week's report...)**
 - !! Oracle problems now resolved !!**
 - **Important that there is regular reporting as previously discussed!**
 - **Come back to this later...**
 - **Incident report attached to agenda of this week's OPS meeting**
 - **<http://indico.cern.ch/conferenceDisplay.py?confId=43733>**
- **Cooling problem at CERN – already reported last week!**
 - **<https://twiki.cern.ch/twiki/bin/view/FIOgroup/PostMortem20081125>**

ASGC – Summary of Main Events (1/2)

CASTOR and SRM services were down from Oct 25 to Nov25

- 25/10/08, Castor frontend throughput observed decreasing to 0% and team ticket raised (ggus ticket open 42913 42878 and 42766).
- 29/10/08, While tracking the ora error, the share memory processes have been forced killed and result in corrupted data files
- 2/11/08, recover the db via RMAN and offline drop data file #25 due to the unrecoverable error. All instances and db services able to startup normally after Nov 2nd 18:00

ASGC – 2/2

- 6/11/08, by dropping one of the table space from dba_rollback_segs specific for one of the instance serving stager service, one of the rollback segments have status NEEDS RECOVERY , and have recreate the undo log and restart everything. Afterward, the stager service had been recovered after redo the statistics.
- 6/11/08 to 25/11/08
 - a) Remove problematic procedures found from error logs, and recompiled invalid procedures.
 - b) Found many index blocks (File #10) corrupted
 - c) Cleanup and enlarge partitions for recovery, and remove obsolete archive logs
 - d) Recreate REDO log files with verification for the error found
 - e) Session number increased to resolve the error - Oracle Code: 12520 ORA-12520: TNS: listener could not find available handler for requested type of server
- ☺ 25/11/08, all Castor, SRM, and Oracle services resumed and verified by experiment data transmission efficiency to almost 100%.

CERN Cooling Problems

- Due to a cooling problem the batch system was shutdown. Submission of new jobs is still accepted but the jobs won't be started. A fix is expected by the end of the day.
 - The cooling of the Computer Centre failed this morning around 08:00 due to a major leak in some cooling pipe. In order to limit the temperature increase in the machine room, the batch system was quickly stopped. The ventilation units were also restarted as soon as possible by the cooling expert in order to cool down the room with fresh air only. The batch system will likely remain down for a few hours, until the cooling system is properly fixed.
 - The leak in the cooling installation is fixed (11:40). The cooling system is being restarted but will only be fully operational at the end of the afternoon. Batch nodes are going to be progressively restarted in accordance with the available cooling capacity and with the importance of the nodes (nodes serving important queues will be restarted first).

GGUS Summary

VO	Alarm	Team	Total
ALICE	0	0	0 (was 0)
ATLAS	1 (test)	4 (was 9)	46 (last week 44)
CMS	0	0	6 (was 7)
LHCb	0	2 (was 1)	7 (was 15)

Baseline Versions Update

- FTS 2.1 (version for SL4)
 - gLite 3.1 SL4 patch 2643 planned for Nov 27th
- Baseline versions linked from daily operations meeting & updated as required

Update on Problems

- We agreed(?) at the last MB F2F that a site had 24h in their time zone (working week) to respond to major problems and try to fix them
 - **“One day’s grace”**
- **A report / update is expected at the daily WLCG operations meeting after this 1 day (problem resolution)**
- **For on-going problems, regular updates highly desirable**
- **Any change of state should be reported**
- **By mail or by updating wiki directly!**
- It really isn't so hard...

A Good Example - PIC

- Thursday last week in the LHCb report there was a comment on the rates obtained at PIC (pre-staging)
 - LHCb (Roberto) - preparation of staging exercise. [RAL](#) & CERN still behaving fine. 20MB/s - most likely will increase. PIC was at 20MB/s. Starting to accumulate big backlog over past 12h - rate has been dropped to 10MB/s.
- Following day Gonzalo joined the call having investigated and gave the following update:
 - PIC reduced to 10MB/s yesterday - Not fully OK (Gonzalo - pre-stage tests needed to reduce because of backlog, believe due to a couple of incidents yesterday and day before. 1 tape drive got stuck with LHCb tape. Caused rates to be reduced. 1 of 2 disk pools for this had an intervention 2 days ago. Caused a reduction in rate for some hours. In principle since yesterday morning both should be ok and hence backlog should clear. Should be able to cope with high rate.)
- Thanks a lot 😊

Storage-Oriented Workshops

- CASTOR F2F meeting at RAL (tbc tomorrow) Feb 18-19
- dCache workshop at FZK Jan 14-15
 - <http://indico.cern.ch/conferenceDisplay.py?confId=45966>
 - Technical exchange on improving stability and reliability of the dCache data management system at WLCG T1 centers
 - 09:00->12:30 T1 Administrators session: Configuration, databases, backups, tools (FTU)
 - 14:00->18:00 T1 Administrators session: HSM, Logging, Problem solving
 - 09:00->18:00 General Session
- Was proposed that these workshops have overlapping agendas with the common goal of getting storage services at the main WLCG sites into good (sustainable) operational status as early as possible in 2009
- **IMHO if we don't do this it will be a missed opportunity...**
- Working towards this goal with Jan & Jos...

Summary

- WLCG Operations Page:
 - <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOperationsWeb>
 - Linked directly from WLCG home page
- WLCG Operations mailing list:
 - wlcg-operations@cern.ch
- WLCG “Service Coordinator on Duty”
 - wlcg-scod@cern.ch
- WLCG Operations review
 - Some spill-over from recent “Data Taking Readiness” workshop
 - Main theme of pre-CHEP workshop?? I reckon so...