WLCG Weekly Service Report

Jamie.Shiers@cern.ch

~ ~ ~

WLCG Management Board, 24th June 2008

https://twiki.cern.ch/twiki/bin/view/LCG/WLCGDailyMeetingsWeek080616

Storage Versions – Present & Future

Component	Version	Comments
CASTOR core	2.1.7-10	will be released this week Tier1s are recommended to upgrade faranno l'upgrade verso meta' Luglio

	2.1.8	will be released the first week of August
		    - Tier0 will upgrade before the end of August
		    - Tier1 will follow

CASTOR SRM	1.3-27 on SLC3	2.7-1 on SLC4 as soon as released
dCache	1.8.0-15p6	fixes a bug with caching credential produced through grid-proxy-init
	1.8.0-15p7	is about to come out. It fixes a problem with checksum verification when copy a file in push mode between 2 dCache sites
StoRM	1.3.20 on SLC4	
DPM	1.6.10 on SLC4	

Power & Cooling

Site	Comments
IN2P3	Had a serious problem this w/e with A/C. Had to stop about 300 WNs - waiting for action this week to repair A/C machine. Keep info posted on website.
INFN	CNAF - suffered serious problem. UPS too heavy & floor collapsed!

- IMHO, a "light-weight" post-mortem, as prepared by various sites for events during the May run of CCRC'08, should be circulated for both of these cases.
- I believe that this was in fact agreed during the February run of CCRC'08
- **I think we should assume that power & cooling problems are part of life and plan accordingly**

Other Site Issues / Questions

Site	Issue
RAL	information gathering plug-ins timing out? -> 0 published as number of running jobs hence not attractive - under investigation (Friday) Rogue CMS user submitted ~10K jobs overloaded batch system - jobs killed, AOK.
CERN	CERN availability currently at 0 (zero) (gstat tests). No published tests since yesterday. Last tests failed. Publishing problem from gstat to SAM(?) Being investigated (Taipei). https://gus.fzk.de/pages/ticket_details.php?ticket=37710
GRIF	Raised question of (possible?) additional ATLAS space tokens at their Tier2 (a suivre)

Databases

- Upgrade to 10.2.0.4 for WLCG cluster & PDB cluster (ALICE + HARP + ...) Interventions < 3.5 hours.
  - Some problems with GridView - some connections fail. Experts investigating...
  - RAC6 - Ethernet switches on critical power (Short interventions on CMSR & LHCBR of a few minutes) - as scheduled.
  - EGEE broadcast for affected services?

- Moved shared SRM to separate RACs (several interventions)

- In order to isolate the Castor SRM service from potential problems on the Castor stager backend (in the case where the stager backend has too many requests and becomes slow to respond), we would like to make a configuration change to all SRM services (Thursday) at 10.30 (just after the dteam/ops SRM intervention). The change will be to lower the stager timeout – this should prevent exhaustion of the SRM thread pool that we observed during CCRC.

Monitoring

- GridView problem of [Monday] - some instability of service wrt Oracle cluster-ware. DBAs recreated service definition - situation looks better - users should not see any instability? Opened service request with Oracle. Incredibly long transactions? Miguel -> developers -> schema review.

- Monitoring / dashboard report: Two issues after the GridView upgrade
  1. the TNS was slightly out of date so the services had to be restarted
  2. a query has stopped working and will be fixed tomorrow so there will be a few hours when the summary data for availability will not be updated.

- CMS dashboard – performance / load problems. Working with DBAs to identify slow / heavy queries. A specific query felt to have been the cause of some such problems is now using an index and runs much faster. A review has been proposed by the DB team and should preferably be scheduled asap (<< data!)

Other Service Issues

- VOMS service issues:
  - VOMS follow-up: Yes FIO don't have a procedure since it is basically impossible to measure when the service is not available since there is no monitoring agent that is a member of the VO. And it seems to be hard to work out externally with out being in the VO if it is working. *Steve*

- \# threads for LHCb LFC increased to 60 taking advantage of DB intervention

- LFC - VOBOXes of LHCb - to be recognized as host certificate by the LFC, a VOBOX DN must finish with **CN=myhostname** or **CN=host/myhostname**. (So, the host certificate - just renewed - for VOBOX for LHCb at NIKHEF is wrong.)

Experiment Issues (1/2)

- CMS – fairly quiet week. Transfer problems to FZK reported Tuesday, CMS SAM SE testing (CASTOR@CERN) failing with write permission denied (Friday), summary of CRUZET-2 (ended last Sunday) reported Wednesday. (100% job success at Tier0, some issues in RECO step – unsuppressed zeroes in ECAL(?) giving very large RECOed files. iCSA08 production finished for all workflows, cleaning up un-needed datasets. HI production runs at Tier1s. Some srm-cms 'misbehaviour' seen well before announced intervention time. All problems disappeared afterwards...

- ALICE – only occasional FTS testing – not enough Tier1 disk space to fully stress sending custodial raw during commissioning exercise. First pass ESD at CERN for transmission to Tier1s for eventual reprocessing

Experiment Issues (2/2)

- LHCb – restart CCRC activities with 'intermediate' version of DIRAC. RAL & IN2P3 to support dcap instead of gsidcap. Problems with info published by RAL (see earlier). More info in elog LHCb observation

- ATLAS – NDGF migrating from RLS to LFC. (Wed). Clarification of naming scheme to be used from now on – data deleted after 8 days:
  - Naming convention for the ATLAS Functional Test datasets
  - ccrc08_run2.0XXYYY.physics_blabla
  - where XX is the week number, this week (started Monday at 0:00 am) is 25, and YYY is sequential.
  - This is just to give the information that we are going, starting tomorrow, to delete the ccrc08_run2.024yyy.blabla data. For the sites this operation is transparent.

WLCG BBQ

- Wednesday June 25th at 17:30 – late at BBQ area

- Wiki page (sign-up, contributions(!)):
  - https://twiki.cern.ch/twiki/bin/view/LCG/WLCGSummerBBQ