



WLCG Service Report

Jamie.Shiers@cern.ch


~ ~ ~

WLCG Management Board, 16th September 2008

[Minutes from daily meetings](#)

Highlights

LHC First Beam day!

- Service Load & Security Problems
- **Post-Mortems – RAL CASTOR/Oracle outstanding**  **Promised for today!**
- IMHO we really need to be timely and precise with these – e-mails are still circulating referring to “Nilo’s suggested DB config changes”.
 - `_kks_use_mutex_pin=false`
- Preparation for WLCG session during EGEE’08

First Beam Day



g "will we be reach
stion now is "wer



ey) services still
e besides...

- This is also true
handling ... put in



Beam day" photo
displays) really sh
ng to rebalance du



Service Load Problems

- Service load has been seen on a number of occasions, including ATLAS conditions DB at various sites, as well as ATLAS CASTOR instance at CERN ([post-mortem](#), including discussion of GGUS alarm ticket follow-up).
 - 10' for CASTOR expert to be called
 - 10' for intervention to start
 - 18' for problem to be identified
 - < 3 hours total from start of problem to confirmation of resolution
- These problems are likely to persist for at least weeks (months?) – we should understand what are the usage patterns that cause them, as well as why they were not included in CCRC'08 tests

Post-Mortem

- 18:10 - problem started
- 19:34 - **GGUS ALARM TICKET** submitted by ATLAS shifter:
- 19:35 - mail received by CERN Computer Centre operator
 - From: GGUS [mailto:helpdesk@ggus.org] Sent: mercredi 10 septembre 2008 19:35 To: atlas-operator-alarm; support@ggus.org Subject: GGUS-Ticket-ID: #40726 ALARM CH-CERN Problems exporting ATLAS data from CASTOR@CERN
- 19:45 - CASTOR expert called
- 19:55 - CASTOR expert starts investigating
- 20:13 - CASTOR expert identifies that the problem is due to a hotspot. The resolution is applied (see below) and ATLAS informed.
- 20:47 - CASTOR re-enabled the diskserver after having confirmed that the requests were better loadbalanced over all servers in pool.
- 20:57 - ATLAS confirms that situation is back to normal

Network Problems

- BNL still “in the headlines” regarding network-related problems. Primary OPN link failed Thursday night when a fibre bundle was cut on Long Island. Manual failover to secondary link. Need to automate such failovers plus continue to follow-up on (relatively) high rate of problems seen with this link
- Network problem at CERN Monday caused 3.5 hour degradation affecting CASTOR

Database Service Enhancements

- Support (on best effort) for CMS and LHCb online databases added to the service team responsibility
- Oracle Data Guard stand-by databases in production for all the LHC experiments production databases (using hardware going out of warranty by the end of the year). Additional protection against
 - human errors
 - disaster recoveries
 - security attacks

WLCG Sessions during EGEE'08

- The idea is to have a panel / discussion session with 3 main themes:
 1. Lessons learned from the formal CCRC'08 exercise and from production activities
 2. Immediate needs and short-term goals: the LHC startup, first data taking and (re-)processing of 2008 data;
 3. Preparation for 2009, including the CCRC'09 planning workshop.
- In each case the topic will be introduced with a few issues followed by a wider discussion involving people also from the floor. Not looking for 'official statements' – the opinions and experience of all are valid and important.
- These panels have worked well at previous events (WLCG workshops, GridPP, INFN etc.) and do not require extensive preparation. It is probably useful to write down a few key points / issues in a slide or two (not a formal presentation!)
- It is also an opportunity to focus on some of the important issues that maybe have not been fully discussed in previous events.

Post-Mortems

- We are now pretty good at preparing timely and detailed post-mortems
- But what happens next?
- e.g. both CASTOR/ATLAS and CASTOR/RAL “post-mort es ” propose actions and other follow-up
- Without inventing a complex procedure, how do we ensure that this happens?
- A: add to MB action list? WLCG operations?

Conclusions

- ☺ **The service is running; problems are responded to and resolved within the time windows that we have established as realistic**
- ☹ Further service hardening clearly required – this will proceed in parallel to the on-going LHC commissioning
 - Consistent follow-up on post-mortem actions required
 - 👉 AFAIK, the interest of “the world” in these activities is unprecedented!