



# WLCG Service Report

**[Jamie.Shiers@cern.ch](mailto:Jamie.Shiers@cern.ch)**

~ ~ ~

**WLCG Management Board, 13<sup>th</sup> January 2009**

# Overview

- The last service report was given on December 16<sup>th</sup> 2008 – one month ago
- Since that time the service has run reasonably – even commendably – well
- Experiments are preparing “post-mortem” analyses of their experiences over Christmas and there is a slot on this at the GDB tomorrow
- We created WLCG operations pages (wiki) for the weeks over Christmas and New Year and people were encouraged to add comments / incident reports directly
- But noone did...

# GGUS Summary – Last Week

VO	USER	TEAM	ALARM	TOTAL
ALICE	<b>1</b>	0	0	1
ATLAS	26	5	0	31
CMS	3	0	0	3
LHCb	2	0	<b>1</b>	3

- Tickets for ALICE are very rare → detail
- We always look at alarm tickets → LHCb

# J "ALICE test attributes VO error to site" T

---

Detailed description: [https://gus.fzk.de/ws/ticket\\_info.php?ticket=45061](https://gus.fzk.de/ws/ticket_info.php?ticket=45061)

The experiment-specific SAM tests for ALICE are failing as such:

Event: Abort

- Arrived = Tue Jan 6 19:27:04 2009 CET
- Host = wms103.cern.ch
- Reason = X509 proxy not found or I/O error
- Source = WorkloadManager
- Src instance = 24407
- Timestamp = Tue Jan 6 19:27:02 2009 CET

However this is reflected in a "failure" of the SAM test and is erroneously attributed to the site.

Note as well that according to SAM, this test has run only four times in the past five days.

Used commandline: looked at our sites Nagios setup

Received error message: see above

Problem affects the whole VO: alice

# LHCb Alarm Ticket (45112)

- Detailed description:  
Stuart Paterson (and Joel Closier as well) claims to have problems registering files in LFC despite using the right FQANs (Role=production). The problem is two-fold. First he can't register file, second LFC does not report any failure (using GFAL Python APIs). I tried to reproduce the problem using lcg-utils command lines from lxplus but I'm not able to reproduce it. I verified that accordingly the ACLs defined in the interested directory, the behavior of LFC is correct.  
Following the extents of the Stuart's problem.  
...

# LHCb Alarm Timeline


santinel	2009-01-08	<b>16:20</b>	assigned (ROC_CERN) Problem affects the whole site: CH-CERN Problem affects the whole VO: lhcb Sent ALARM mail to mail address lhcb-operator-alarm@cern.ch
aretico	2009-01-08	16:31	in progress (ROC_CERN)
Computer Operations	2009-01-08	16:47	Added attachment mailbody.Thu._8_Jan_2009_17.45.47_.0100.txt
Computer Operations	2009-01-08	<b>16:47</b>	Public Diary: For your information, The Data Operations Piquet has been called. Please standby. Regards,Pascal Sicault operator on duty
Uli	2009-01-08	16:53	Checked with LFC expert who already started to investigate this. Copy&paste from offline conversation: ----8<---- Hi Roberto, The only thing I can see in the LFC logs is the following:
Slemaitr	2009-01-09	16:05	Solved (ROC_CERN)

LHCb ( Roberto) 1. issue reported last week with LFC was DIRAC handling of registration. First checks if GUID is registered if not tries to register. If yes "gives up" - problem on client side.

# Service Incidents Requiring Follow-up

A user running a test on Wednesday night triggered several deadlocks in the Name Server SQL code that

The problems affecting all CASTOR instances on Wednesday evening have been understood as being caused by a bug in



■ eth0 in aver:920.2M max:2.0G min:442.9M curr:1.2G  
■ eth0 out aver:5.0G max:8.6G min:762.0M curr:5.0G  
■ eth1 in aver:1.2k max:5.3k min:282.8 curr:299.8  
■ eth1 out aver:489.6 max:791.3 min:436.4 curr:445.9

# Other Issues

- Issues related to shared s/w area have come up a couple of times recently
- This need to be taken up and addressed at an appropriate forum, e.g.
- HEPiX – techniques for running large scale shared filesystems
- s/w installation and other tools – avoid use of hard-coded absolute paths
- AF – compiler version(s) required for SL5



# Outlook

- The amount of (valuable) information reported in the daily operations meetings continues at a high level – the calls can last up to 30' these days
  - How will this “scale” to LHC running?
- It is important that we continue to follow-up on significant service degradations and incidents
  - Services that are significantly degraded and / or for a long period of time
- There is increasing evidence that more “best practices” knowledge sharing would help many sites and hence the overall service
  - More cooperation with HEPiX on specific issues?
- How do we ensure adequate cross-reporting between site internal meetings and “the Grid”?
  - AFAIK I only have access to those from CERN & GridPP...

# Reminder

- Late last year I suggested a target whereby the weekly service report would essentially be “business as usual” 3 weeks out of 4
- We are still a long way from that target...