



WLCG Service Report

Maria.Girone@cern.ch, Harry.Renshall@cern.ch

~ ~ ~

WLCG Management Board, 7th December 2010


Covers 2 week period from 22 November

Introduction

- Generally smooth operation on experiment and service side
 - 4 GGUS alarms
- Service Incident Reports received:
 - 16/11 and 26/11: GGUS Database down
 - 29/11: CERN CASTOR name server
- Service Incident Reports to be produced / updated:
 - HI run: IN2P3 ATLAS slow transfers
 - 27/11, 30/11, 02/12, 04/12: CERN ATLAS Offline DB node reboots
 - Coincidence with log rotate – now turned off – still to be understood
 - RAL disk server problem – still ongoing – but interim SIR
- Some site availability issues – commented directly in Availability Plots – in particular several sites suffered **power failures**
 - **FNAL 22/11, NDGF 24/11, RAL 01/12**

GGUS summary (2 weeks)

VO	User	Team	Alarm	Total
ALICE	2	0	2	4
ATLAS	15	131	5	151
CMS	4	3	2	9
LHCb	1	18	1	20
Totals	22	152	10	184



ATLAS ALARM->CERN-CNAF TRANSFER FAILURES

What time UTC	What happened
2010/11/24 9:00	GGUS ALARM ticket opened, automatic email notification to t1-alarms@cnafe.infn.it AND automatic assignment to ROC_Italy.
2010/11/24 9:04	Submitter links to TEAM ticket GGUS:64571 for details and escalates ALARM already!!??
2010/11/24 9:14	Site starts investigation. Atlas blacklists DATATAPE use at CNAF due to too many errors.
2010/11/24 11:19	'Solved'. Maybe due to a misconfiguration for the VOMS servers (BNL voms reference has disappeared from the end-point). Submitter sets status 'verified'

- https://gus.fzk.de/ws/ticket_info.php?ticket=64577

ALICE ALARM-> FZK CREAM-CE

What time UTC	What happened
2010/12/05 11:32 SUNDAY	GGUS ALARM ticket opened, automatic email notification to de-kit-alarm@scc.kit.edu AND automatic assignment to NGI_DE.
2010/12/05 13:32	Site acknowledges ticket. Segmentation fault since midnight. Maybe due to a corrupted file system. Needs to wait for Monday to fully fix.
2010/12/06 8:20	/proc unreadable on many compute nodes.
2010/12/06 11:58	'solved' without needing a full PBS reset, as originally feared. Submitter 'verified' at 13:29

- https://gus.fzk.de/ws/ticket_info.php?ticket=65015

ATLAS ALARM-> SARA

MATRIX LFC REG.FAILURES

What time UTC	What happened
2010/12/05 19:18 SUNDAY	GGUS ALARM ticket opened, automatic email notification to grid.support@sara.nl AND automatic assignment to NGI_NL.
2010/12/06 8:40	Site acknowledges ticket.
2010/12/06 10:10	Set to 'solved' by the site by extending the table space.
2010/12/06 11:38	Re-opened by the submitters due to new errors seen in DDM. Question at the WLCG daily meeting was why no action was taken in the evening of 5/12.

- https://gus.fzk.de/ws/ticket_info.php?ticket=65019

CMS ALARM->CERN CASTOR

What time UTC	What happened
2010/11/29 14:25 SUNDAY	GGUS ALARM ticket opened, automatic email notification to cms-operator-alarm@cern.ch AND automatic assignment to ROC_CERN.
2010/11/29 14:41	Site operator acknowledges ticket. Piquet called.
2010/12/03 11:21	Set to 'solved' by the site as a known issue but without diagnosis. Remains 'in progress' in GGUS, following a re-classification in Remedy PRMS. To be followed up.

- https://gus.fzk.de/ws/ticket_info.php?ticket=64726
- **SIR:** CERN CASTOR name server

GGUS SIRs



- Down 3.5h on 16/11 and 1.5h on 26/11
 - No web access / no ticket update in both cases
1. 16/11: outage isn't understood up to now – more details in SIR – maybe load related
 1. 26/11: the reason for the outage was an unscheduled reboot of the Oracle DB cluster failing for all nodes. The DB reboot was caused by a combination of hardware failure and miss-configuration. The hardware failure made the DB cluster lose some essential control files. The miss-configuration prevented the cluster reboot.

CASTOR NAMESERVER



- Problem with the Name Server database. The root cause seems to be related to a problem accessing one of the db datafiles. The database set the problematic datafile offline.
- All CASTOR instances were affected, the service was unavailable between 15h and 16h Geneva time (UTC + 1)

When	What
29-Nov-10 15:04:00	Probe tests start failing
29-Nov-10 15:07:00	Problem diagnosed as name server problem: nsis -l / returns 'Internal error'
29-Nov-10 15:07:00	Oracle Support contacted
29-Nov-10 15:27:00	Mail to info experiments. Downtime entered in GOCDB. Incident entered in SSB.
29-Nov-10 15:28:00	Alarm ticket from CMS
29-Nov-10 15:40:00	Database recovered
29-Nov-10 15:50:00	Service restored
29-Nov-10 16:00:00	Mail to info experiments. Incident updated in SSB.



ATLAS
7 Days from 2010-11-22 to 2010-11-29





ALICE
7 Days from 2010-11-22 to 2010-11-29



CMS
7 Days from 2010-11-22 to 2010-11-29



LHCb
7 Days from 2010-11-22 to 2010-11-29



Analysis of the availability plots

ATLAS

1.1 INFN: SRM tests were failing: lcg-cp (no files have been lcg-cp on this endpoint), lcg-cr (error reading token data header: connection closed), lcg-del (unary operator expected) possibly connected to INFN transfer and storage problems (original problem reported on Friday the 19th – Network Problems GGUS: 64459). Site reported that it was a temporary storage outage.

ALICE

2.1 FZK: VOBox User Proxy Registration test was failing (the user is not allowed to register his proxy within the VOBOX).

2.2 Green box for IN2P3 & FZK: CE tests were failing. SAM CE tests should be ignored as ALICE is using CREAM CE. There are no CREAM CE direct submission tests for ALICE yet (work in progress).

CMS

3.1 ASGC: The CMS software installation SAM test was failing in 5 CEs of ASGC. CEs started publishing CMS SW releases that were not yet installed - Savannah: 117339.

LHCb



4.1 IN2P3: Issues with IN2P3 shared area – GGUS:59880 (performance issues) and GGUS:62800 (SW installation). CE-tests were failing: install, job-Boole, job-DaVinci, job-Brunel,job-Gauss, sft-vo-swdir ('took more than 150 seconds while 60 seconds are expected' OR 'software check failed'). SRM DiracUnitTestUser test failing occasionally (failed to put file to storage.).

4.2 PIC: CE tests were failing: sft-vo-swdir (took more than 150 seconds while 60 seconds are expected) and lhcb-job-DaVinci (software not Installed).

4.3 SARA: CE tests were failing: sft-vo-swdir (took more than 150 seconds while 60 seconds are expected). SRM lcg-cr test failed once on the 28th: CGSI-gSOAP running on volhcb29.cern.ch reports Could not load client credentials.



ATLAS

7 Days from 2010-11-29 to 2010-12-06





ALICE

7 Days from 2010-11-29 to 2010-12-06





CMS

7 Days from 2010-11-29 to 2010-12-06



LHCb

7 Days from 2010-11-29 to 2010-12-06



Analysis of the availability plots

ATLAS

1.1 NDGF: SRM tests were timing out on the 30th (a day before the planned downtime of the SRM) while trying to write a file into the TAPE buffer disk. GGUS:64843.

ALICE

2.1 Green box for IN2P3, NIKHEF and RAL: CE tests were failing. SAM CE tests should be ignored as ALICE is using CREAM CE. There are no CREAM CE direct submission tests for ALICE yet (work in progress).

CMS

3.1 RAL: The CE sft-job test was failing with the following proxy-related errors: 'Got a job held event, reason: Globus error 131: the user proxy expired (job is still running)' & 'Failed to get expiration time of proxy'. The failing test was running as user 'cmssgm', which is only allowed to have one job at a time. One job had been stuck for a few days, preventing other cmssgm jobs from running. Savannah:118201.

LHCb

4.1 IN2P3: Minor issues with IN2P3 shared area – GGUS:59880 (performance issues): The sft-vo-swdir CE test was timing out (took more than 150 seconds while 60 seconds are expected).

Other Issues

- Shared s/w area: very common – and repetitive – cause of problems to experiments
- Some instabilities seen by ATLAS at IN2P3 during HI run – no longer seen
 - Good communication established between site and experiment (daily reports) which will be useful also at daily operations call
 - SIR in preparation

Summary

- HI data processing and proton data reprocessing proceeding smoothly
- Data taking for 2010 now over