



WLCG Service Report

Andrea.Valassi@cern.ch

~ ~ ~

WLCG Management Board, 23rd November 2010

Introduction


- Generally smooth operation on experiment and service side
 - Coped well with higher data rates during the HI run (CMS from CASTOR: 5 GB/s)
- One Service Incident Report received:
 - IN2P3 shared area problems for LHCb (interim [SIR](#) – [GGUS:59880](#))
- Two more SIRs are pending:
 - CASTOR/xrootd problems for LHCb at CERN ([GGUS:64166](#))
 - GGUS unavailability on Tuesday November 17th
- Three GGUS ALARMS
 - CASTOR/xrootd problems for LHCb at CERN ([GGUS:64166](#))
 - ATLAS transfers to/from RAL ([GGUS:64228](#))
 - CNAF network problems affecting ATLAS DDM ([GGUS:64459](#))
- Other notable issues reported at the daily meetings
 - *Slow transfers to Lyon for ATLAS (dcache problems)*
 - *Many top-priority tickets are open (e.g. [GGUS:63631](#), [GGUS:64151](#), [GGUS:64202](#))*
 - Security updates in progress (CVE-2010-4170)
 - BDII timeouts for ATLAS at BNL due to OSG network problems ([GGUS:64039](#))
 - Database problems for ATLAS Panda and PVSS at CERN (no GGUS ticket)

IN2P3 (intermediate) SIR

- Since some months LHCb is suffering from problems related to the IN2P3 shared area on AFS:
 - up to 6% of all jobs fail due to timeout during software setup (GGUS:59880)
 - software installation jobs fail (GGUS:62800)
- ATLAS had seen a similar problem with software setup timeouts
 - workaround for ATLAS (increase timeout) recommended also for LHCb
- Follow-up is still in progress
 - Separate WNs for LHCb (where ATLAS jobs are excluded) are being deployed
 - Tuning and tests are ongoing

GGUS summary (2 weeks)

VO	User	Team	Alarm	Total
ALICE	2	0	0	2
ATLAS	20	104	2	126
CMS	5	4	0	9
LHCb	2	14	1	17
Totals	29	122	3	154



LHCB ALARM->CERN CASTOR ACCESS FAILURE

What time UTC	What happened
2010/11/11 8:43	GGUS ALARM ticket opened, automatic email notification to lhcb-operator-alarm@cern.ch AND automatic assignment to ROC_CERN.
2010/11/11 8:58	Site acknowledges ticket. CASTOR piquet called.
2010/11/11 10:32	'Solved' by explaining that incoming proxies with VOMS Role=NULL are not recognised as valid members of the root group /lhcb.
2010/11/11 10:54	LHCb supporter requests to add this FQAN in the mapping of the xrootd redirector.

- https://gus.fzk.de/ws/ticket_info.php?ticket=64166

ATLAS ALARM->CERN-RAL TRANSFER FAILURES

What time UTC	What happened
2010/11/14 7:11 SUNDAY	GGUS ALARM ticket opened, automatic email notification to lcg-alarm@gridpp.rl.ac.uk AND automatic assignment to ROC_UK/Ireland.
2010/11/14 7:49	Site acknowledges ticket. CASTOR expert investigating.
2010/11/14 9:19	SRM & LSF services restarted at RAL. Submitter reports persistent problems for the rest of the day. Links to TEAM ticket GGUS:64224 for details.
2010/11/15 11:19	'Solved' by reducing the allowed number of jobs on the batch farms. Submitter sets status 'verified

- https://gus.fzk.de/ws/ticket_info.php?ticket=64228

ATLAS ALARM->CERN CNAF TRANSFER FAILURES

What time UTC	What happened
2010/11/20 20:32 SATURDAY	GGUS ALARM ticket opened, automatic email notification to t1-alarms@cnafe.infn.it AND automatic assignment to ROC_Italy.
2010/11/20 20:53	Site acknowledges ticket. Investigation till midnight gave no understanding of the problem reasons.
2010/11/21 10:45	Downtime recorded in GOCDB. Atlas blacklists the site.
2010/11/22 14:17	SRM restarted, CERN jobs accepted again. Site contacts suggest to close the ticket, even if the problem is not explained.



- https://gus.fzk.de/ws/ticket_info.php?ticket=64459

Support-related events since last MB

- GGUS TEAM and ALARM tickets will be equipped with a 'Problem Type' field, like USER tickets, in order for periodic reporting to better show weak areas in support.
- Individual tickets were opened against each middleware-related GGUS Support Unit to inform them about assignment ONLY by the DMSU. Some middleware experts refuse to be hidden behind the 2nd level EGI body DMSU in GGUS. This affects implementation of the new workflow in ticket assignment, at the next release (Dec. 1st).
- The suggestion to create TEAM instead of USER tickets via the CMS-to-GGUS bridge was rejected.
- The request to be able to convert TEAM tickets into ALARMS was accepted.
- There were 3 true ALARM tickets since the Nov. 9th MB (2 weeks), Details in the previous slides.



ATLAS

7 Days from 2010-11-08 to 2010-11-15





ALICE

7 Days from 2010-11-08 to 2010-11-15




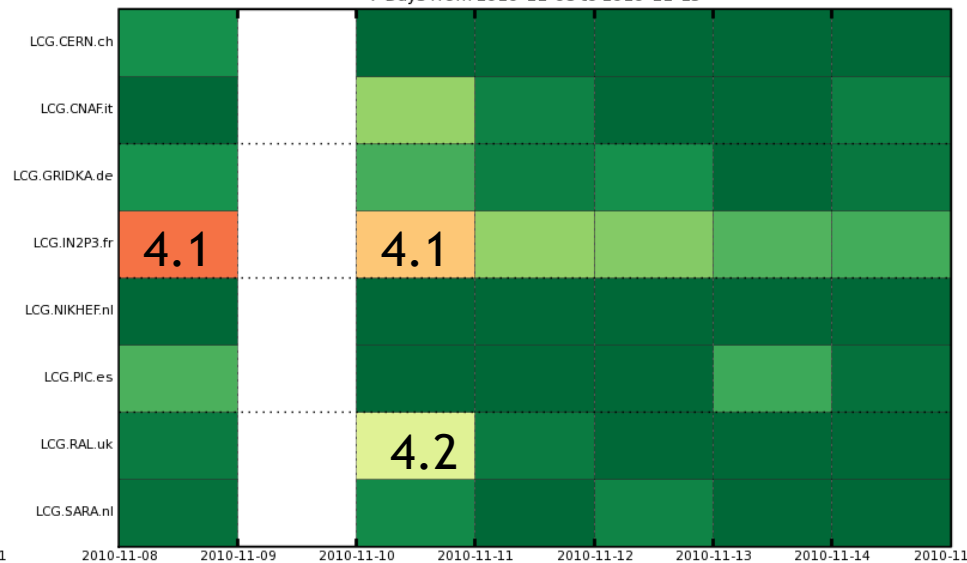
CMS

7 Days from 2010-11-08 to 2010-11-15



LHCb

7 Days from 2010-11-08 to 2010-11-15



Analysis of the availability plots

Tuesday 9th of November: We are still experiencing problems with missing data for one day of the availability report (week 101108) as reported on the daily WLCG operations meetings on the 10th on November.

ATLAS

Green box for BNL: Problems with SAM BDII - availability info is false as tests were failing due to issues not to do with the site.

1.1 IN2P3: SRM tests were timing out after 600 seconds due to a disk space issue + SRM highly loaded by SRMPUT - GGUS:64164 and GGUS:64151.

1.2 RAL: SRM tests were failing: lcg-cp (timing out after 600 seconds), lcg-cr (timing out after 600 seconds) and lcg-del (unary operator expected) – GGUS:64228.

ALICE

2.1 Green box for SARA: 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 FNAL: SE type SRMv2 not published - After a SAM-SRM error at T1_US_FNAL, no further SAM tests were running - GGUS:64084.


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'), condDB (problems opening the database – error found in the output job) and sftp-job (globus error 126: it is unknown if the job was submitted). SRM tests failing: lcg-cr (connection time out, no accessible BDII), diracUnitTest (assertion error).

4.2 RAL: Unscheduled outage: service at risk while SRMs upgraded by rolling change 11th Nov 10h until 11th Nov 14h. SRM tests where failing: lcg-cr, diracUnitTest and lhcb-fileaccess (could not open connection to srm-lhcb).


ATLAS

7 Days from 2010-11-15 to 2010-11-22




ALICE

7 Days from 2010-11-15 to 2010-11-22



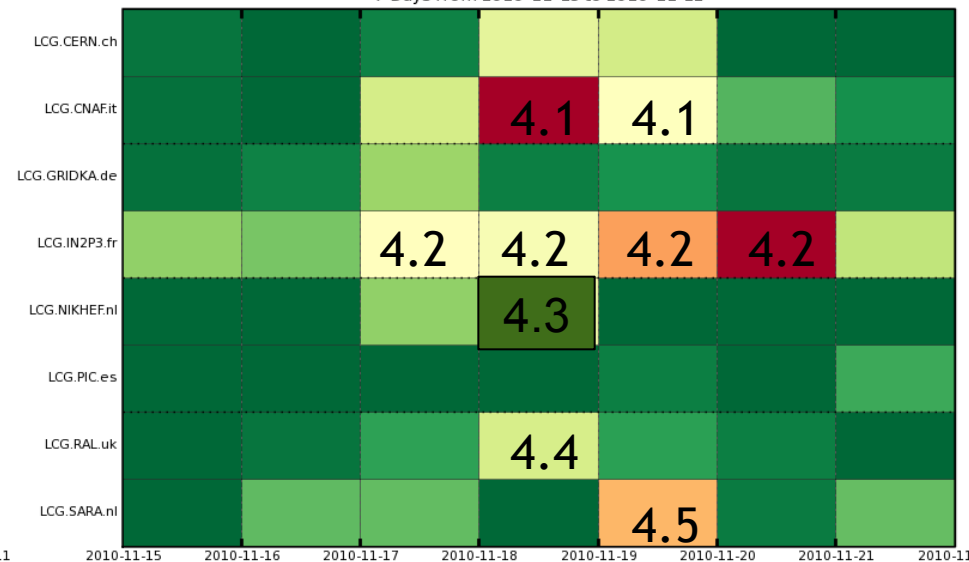
CMS

7 Days from 2010-11-15 to 2010-11-22



LHCb

7 Days from 2010-11-15 to 2010-11-22



Analysis of the availability plots

ATLAS

1.1 INFN: SRM tests were failing as reported on Friday the 19th: 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 due to INFN experiencing network problems (GGUS: 64459).

ALICE

2.1 Green box for SARA: 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 KIT: SRM tests were failing as reported on the 16th: lcg-cp (timeout after 1800 seconds), lcg-ls, lcg-gt and lcg-gt-rm-gt (globus_ftp_client: server responded with an error). Probably an indication that CMS skimming was overloading the system.

3.2 CNAF: CE and SRM tests were failing. CE: cms-analysis (open failed with system error 'Stale NFS file handle'), sft-job (Got a job held event, reason: Globus error 25: the job manager detected an invalid script status). SRM: lcg-cp, lcg-gt, lcg-ls & lcg-gt-rm-gt (Client transport failed to execute the RPC) and lcg-ls-dir (time out after 1800 seconds). Possibly due to INFN experiencing network problems (GGUS: 64459 & 64462).

3.3 RAL: Scheduled outage of the SRM from the 16th until the 18th: Upgrade of the CMS Castor instance to version 2.1.9.

3.4 FNAL: SRM tests were failing/timing out due to the heavy usage of the storage element - some of the components responded within larger time intervals. GGUS: 64463.

LHCb

4.1 CNAF: SRM tests were failing: lcg-cr (communication error on send), diracUnitTest (failed to put file to storage) and fileAccess (invalid argument). Possibly due to INFN experiencing network problems (GGUS: 64459).

4.2 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'), condDB (problems opening the database – error found in the output job) and sftp-job (globus error 126: it is unknown if the job was submitted). SRM tests failing: lcg-cr (connection time out, no accessible BDII), diracUnitTest (assertion error).

4.3 Green box for NIKHEF: The lhcb-availability test was failing. The lhcb-availability is not considered a 'critical' test and we are investigating why this tests is incorrectly included on the results of the critical tests.

4.4 RAL: SRM tests were failing (open/create error: read only file system): lcg-cr (the server responded with an error) and diracUnitTest (failed to put file in storage).

4.5 SARA: SRM tests were failing: lcg-cr (connection time out), diracUnitTest (failed to create directory on storage), fileAccess (no accessible BDII).

Conclusions

- Business as usual – busy but successful
- Some issues need follow-up (e.g. ATLAS transfers to Lyon)
- WLCG is meeting the challenges of HI data taking