

Table of Contents

Week of 090413.....	1
WLCG Baseline Versions.....	1
WLCG Service Incident Reports.....	1
GGUS Team / Alarm Tickets during last week.....	1
Weekly VO Summaries of Site Availability.....	1
Daily WLCG Operations Call details.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	1
Wednesday.....	2
Thursday.....	4
Friday.....	4

Week of 090413

WLCG Baseline Versions

WLCG Service Incident Reports

- This section lists WLCG Service Incident Reports from the previous weeks (new or updated only).

GGUS Team / Alarm Tickets during last week

- Tickets: see https://gus.fzk.de/ws/ticket_search.php and select "VO" - this will give all tickets, including team & alarm - or use the following link:
 - ◆ In the GGUS Escalation reports every Monday
- Please observe the procedure to become a LHC Experiment VO TEAM member and other most recent GGUS FAQs.

Weekly VO Summaries of Site Availability

- ALICE
- ATLAS
- CMS
- LHCb

Daily WLCG Operations Call details

To join the call, at 15.00 CE(S)T Monday to Friday inclusive (in CERN 513 R-068) do one of the following:

1. Dial +41227676000 (Main) and enter access code 0119168, or
2. To have the system call you, [click here](#)

General Information

- Web archive: <https://mmm.cern.ch/public/archive-list/w/wlwg-operations>
- CERN IT status board: <http://it-support-servicestatus.web.cern.ch/it-support-servicestatus/>
- elogs: <https://prod-grid-logger.cern.ch/elog/CCRC'08+Logbook/>

See the weekly joint operations meeting minutes

Additional Material:

Monday:

(no meeting today - Easter!)

Tuesday:

Attendance: local(Jean-Philippe, Harry, Roberto, Ewan Sophie, Olof, Alessandro, Steve, Nick, Julia);remote(Jeff, Gareth, Michel, Daniele, Angela, Luca).

Experiments round table:

- ATLAS report [↗](#) - (Alessandro) Two important issues: observed MCDATATAPE buffer at FZK is full. A ticket has been submitted and as a first measure more space was added. ATLAS would like to know more details. Second problem: observed only 55% job efficiency at INFN Tier-1 while other Tier-1s were much higher. Would like to understand why. Roberto wondered if this could be related to the 'shared area' in GPFS? Not directly, it seems. It seems that only certain experts could access the area while others couldn't. Still under investigations.
- CMS reports - (Daniele) some notes from Easter period, which was relatively quite. A SLS problem on April 7th, which took two tickets to be solved. The experiment shifter is relying on the SLS plot for available space CASTORCMS/T0EXPORT and would like to know if the ticket was submitted correctly. Sophie: do you have the ticket number? Daniele will put the Remedy number in his twiki. The ticket should probably have been submitted to GGUS instead from where it would be routed via CERN site to CASTOR support.
- ALICE -
- LHCb reports - (Roberto) Since last week LHCb maintains their operation report in twiki in a similar manner as CMS. Tier-1 issues: PIC, SARA, INFN, FZK WMS submission is failing for LHCb certificate. Michel has seen a similar problem for ALICE certificates and found that it was related to the short CRL for certificates created by the CERN CA. Roberto said that that is likely to be cause. Other issue: wrong locality returned by Lyon SRM/dCache seems to have been fixed and ticket can be closed. A GGUS ticket has been opened to CNAF. Over Easter LHCb ran 15,000 successful jobs. This week the Tier-1s will receive data and jobs for the FEST. Harry reported that LHCb CVS only works partially.

Sites / Services round table:

- GRIF (Michel): 10 days started to test the WMS new version. The tests are successful so far and we have seen improvements in the performance. Only submission to LCG CE tested so far.
- CNAF (Luca): problem experienced with SRM endpoint at CNAF since a few days. All requests from certain users are failing. The problem is being investigated. Last Thursday somebody from LHCb had raised an issue with an old GGUS ticket opened against CNAF. Luca followed this up and found that the ticket had been answered to by the developers. Roberto: this seems to have been a misunderstanding within LHCb.
- FZK (Angela): some problem with CEs and WMS submission being investigated. Root cause not understood yet.
- CERN: starting on Tuesday last week after a glibc upgrade, cron stopped working on some 30+ machines at CERN, notably VOMS and SRM, which in turn caused CRLs to not be updated. A Service Incident Report is being prepared here: [VomsPostMortem2009x04x10](#)

AOB:

Wednesday

Attendance: local(Jean-Philippe, Harry, Olof, Nick, Antonio, Alessandro, Roberto, Patricia, Miguel, Sophie, Julia);remote(Michael, Gareth, Jeff).

Experiments round table:

- ATLAS report [↗](#) - (Alessandro) reprocessing jobs for INFN Tier-1 finally reached 99% of efficiency (50% before). The reason was because the pool accounts for atlprod was exhausted, related to the lcg-expired-gridmapdir. ASGC asked to clean the LFC catalogue. Thereafter they will try with the

Tuesday:

Australian Tier-2 to transfer files, which belongs to ASGC cloud but for this exercise it will be associated to TRIUMF as a test of the resilience to a complete Tier-1 outage.

- CMS reports - apologies from Daniele couldn't attend
- ALICE - (Patricia back from holidays) WMS 3.2 upgrade at GRIF should allow ALICE to go to 50k jobs/day without job-collections. The same test will run against CERN (3.1) in the next hours. Antonio: where has the 3.2 version been taken from (it should still be in certification)? Patricia doesn't know but presumably the version was taken directly from the developers. *After the meeting Michel Jouvin posted a reply to this question at the wlcg-op list: I apologized I was not able to attend to today as I am in holidays until April 27th. But to answer the question asked by Antonio to Patricia about the origin of WMS 3.2 used at GRIF, yes this is the version currently in certification. This is a "private initiative" between developpers and GRIF and this is not intended to replace the certification in any way. This was thought as the easiest way to fix some problems we were experimenting (in particular related to match making performance and problems, supposed to be fixed in 3.2) and get an early feedback from a site with enough expertise and able to do a controlled deployment in a production environment (including ability to quickly rollback). GRIF who is running French WMS (2) for ALICE discussed with ALICE the possibility to upgrade one of their WMS. Thus the report from Patricia.*
- LHCb reports - (Roberto) first important news: since yesterday SAM jobs fail because of a LHCb scripts but it doesn't affect the site availability computation. Thanks to suggestions from Michel Jouvin yesterday, the CRL problem on WMS has now been solved. CVS problem also reported yesterday has been solved (a node as black-hole). MC production is currently running at 5k jobs in the system. **CNAF BDII issue: publishing incorrect information since their shutdown two weeks ago. This causes a lots of trouble for LHCb match-making. This is fairly important given the imminent FEST run.** As nobody from CNAF was connected to the wlcg-op conf call Roberto will contact them directly.

Sites / Services round table:

- BNL (Michael): operational issue due to some ATLAS production 'tasks' concerning the merging of multi-10k files at an average size of 30MB. All files are on tape. Currently processing 55k concurrent file staging requests, with 77k more in the queue. Since the destination for these files from tape are the same pools data replication requests are served from transfer timeouts are observed. As a result the transfer efficiency as reported by the DDM Dashboard can go down as low as 70% but usually is observed around 80%.
- RAL (Garteth): failing SAM tests on one of the CASTOR CMS instances. Looks to have been a site network problem and it was solved in the morning. A planned intervention was successfully performed this morning on one of the ATLAS 3D databases. Alessandro: do you think this could be related to the ticket about lost ATLAS files at RAL? Gareth don't think it is related to the 3D database issue but he will check for the ticket and follow-up ('file has no copy on tape').
- NIKHEF (Jeff): had previously capped the number of ATLAS jobs at NIKHEF because they were using too much bandwidth. The cap has now been removed.
- CERN (Miguel): started a series of small (transparent) changes to CASTOR Tier-0 services for achieving a more redundant deployment setup. The intervention was scheduled in GOCDB as 'at risk' but the CERN-PROD site was flagged as unscheduled downtime so Miguel canceled it and put the intervention on hold in order to understand better the impact of this. Nick will follow it up.

Middleware (Antonio): glite 3.1 update 44 was released yesterday. Affects the CREAM CE and associated yaim parts. The release has been successfully used by the experiments and sites are encouraged to upgrade. Issue with submission to multiple WMS where one is draining: patch 2.9.2.8 should fix this. ICECREAM patch was removed because of delays in certification. See ScmPps for details.

AOB:

Thursday

Attendance: local(Julia, Sophie, Miguel, Olof, Harry, Alessandro, Roberto);remote(Michael, Angela, Gareth, Danielle).

Experiments round table:

- ATLAS report [🔗](#) - Two issues at FZK - there was an 8 hours unscheduled downtime yesterday for ATLAS but which is understood and a repeat of the disk buffer to tape being full as it was 3 days ago when the migration scheduler had stopped. One network issue between CNAF and TRIUMF where the data transfer rate is stuck at 1 MB/sec since a week - under investigation.
- CMS reports - Asian sites are being tested before restarting site operations. ASGC is now coming green nicely while TIFR (India), Beijing and the ASGC Tier2 have had good metrics for weeks. Islamabad is still ramping up. Site readiness will be discussed next week in the CMS meeting in San Diego. Central shift operations at CERN will soon be restarting and will include a bridge between the CMS Savanna problem reporting system and GGUS.
- ALICE -
- LHCb reports - 1) There was another failure of the CERN CVS system last night, the second this week. Escalation of the support process will be needed. 2) A phone call was held with CNAF over their BDII publishing wrong numbers of queued jobs. 3) Many jobs are crashing at Tier 1 sites trying to open the conditions database - thought to be the way Persistency accesses the LFC. 4) at RAL a new Voms host certificate was not properly deployed on worker nodes. LHCb will add a Voms proxies check to their SAM suite.

Sites / Services round table:

- TRIUMF: We've been running at about 30% CPU capacity since the easter long weekend at a reduced heat load in order to match a reduced cooling capacity. There was a slow leak on a threaded fitting that prevented the main cooling pumps to sense enough differential pressure in order to provide adequate refrigerant flow to our inrow coolers. The leak has been fixed. It's minor but requires time to fix because the affected loop needs to be emptied, sealed, vacuum/pressure tested, then refilled. We will be adding heat load gradually and we expect all worker nodes to be in production later today. (Reda)
- CERN (email from E.Grancher): The CASTOR-ORACLE BigID problem that has caused serious several service interruptions has been definitively understood. An in-depth explanation and fix for it is being prepared.

AOB: N.Thackray followed up on yesterdays question as to whether an 'at risk' unscheduled downtime counted against a sites availability and the answer is that it does not. Gareth confirmed that labelling of unscheduled or scheduled is automatically assigned at the time of entry into the GocDB.

Friday

Attendance: local(Harry, Nick, Jean-Philippe, Roberto, Patricia, Olof, Jacek, Simone, Alessandro);remote(Michael, Gonzalo, Jeff).

Experiments round table:

Thursday

- ATLAS reports - 1) The FTS proxy delegation failure happened at IN2P3 but this is the old server code. ATLAS would like to know site plans to upgrade to the new version. 2) There were many overnight request timeouts at RAL - waiting for more information. 3) FZK had troubles keeping up with tape migration of the MCDISK buffer pool and increased it from 1 to 5 TB to clear the backlog. They then returned to 1 TB and the buffer filled up again. There was no reply to an email sent about 16.00 so ATLAS sent an alarm ticket at about 20.00 yesterday. FZK acknowledged the problem but think this was not sufficiently important for such a ticket - to be followed up. 4) There was a degradation of WMS performance yesterday. Sophie reported there had been a stuck job on an ATLAS WMS - fixed this morning.
- CMS reports -
 - ALICE - A disk filled on a central ALICE server resulting in the number of concurrent jobs grid wide dropping from thousands to a few hundred. This has been fixed and job numbers are ramping up again. Two new CREAM-CE sites have been added - Subatech (Nantes) and IHEP (Russia). Prague has added the required second VO box but it has firewall connectivity problems.
 - LHCb reports - 1) The job failures yesterday at Nikhef and IN2P3 are now explained by the pre-pended "root:" string to the returned tURL.(trick to instruct GAUDI application to open files as ROOT files) 2) The problem of jobs crashing accessing the LFC is still under investigation but seems to be that the thread pool in LFC becomes exhausted due the the way Persistency is accessing it. 3) The CNAF BDII publishing wrong job queue numbers is now fixed.

Sites / Services round table:

BNL (ME): There are still over 100000 staging requests from the ATLAS production tasks (pile up and HITS merging) and this causes a high load on the dCache/pnfs server resulting in an unacceptably high failure rate for DDM transfers. BNL is discussing reducing the load with ATLAS production operations. Simone pointed out that in future small hits files will be merged from their disk copies.

PIC (GM): Are seeing very low cpu efficiencies for LHCb jobs. Roberto will see if their Persistency-LFC problems are the cause.

CERN Databases (JC): at 18.30 yesterday the Oracle table space used by LCG-SAM filled up and new data could not be inserted till 9 am today so some test results have been lost.

CERN CASTOR (OB): There was a degradation on CASTORCMS this morning due to the reset of an Oracle execution plan.

AOB:

-- JamieShiers - 09 Apr 2009

This topic: LCG > WLCGDailyMeetingsWeek090413

Topic revision: r21 - 2010-06-11 - PeterJones



Copyright &© 2008-2022 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback