

# Table of Contents

<b>Week of 090316.....</b>	<b>1</b>
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:.....	3
Wednesday.....	4
Thursday.....	5
Friday.....	7

# Week of 090316

## WLCG Baseline Versions

## WLCG Service Incident Reports

- This section lists WLCG Service Incident Reports from the previous weeks (new or updated only).
- March 16 - ATLAS CASTOR and SRM production instances were down for approximately 12 hours: report

## GGUS Team / Alarm Tickets during last week

- Tickets: see [https://gus.fzk.de/ws/ticket\\_search.php](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
  
- Summary

## 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:

Attendance: local(Daniele, Ricardo, Jamie, Harry, Andrea, Nick, Markus, Patricia, Roberto, MariaDZ, Gavin, Alessandro, Maria, Steve);remote(Angela, Gareth, JT).

Experiments round table:

- ATLAS (Alessandro) - Friday night observed problem with castor at CERN - thanks to people who worked hard to fix problem! Small issue - information in GOCDDB has been added almost "on solution" of the problem - clearly one of the last things you do but we try to have automatic procedure if site -> unsched down can take out of production automatically. -> very useful if in GOCDDB. Ricardo - have changed procedure to address this. FZK: reincluded in func. tests Saturday night and in "full production" Sunday morning. One problematic channel -> Uni. Munich. Lyon observed FTS error over w/e - SRM load related? AOD replication restarted Friday: BNL, RAL & IN2P3. For BNL not clear why data is not going there .. under investigation. Andrea - problem with CASTOR? Ricardo - DB problem. Brian - in UK see issues on a couple of channels. Auto-reduce FTS settings on channels.
- CMS full daily report (Daniele) - top 3 site-lrelated comments: 1) file not going to CASTOR tape since > 10 days - may become a problem when we need to access these files. 2) stagein issues at FZK - following up in Sav. & GGUS. CMS contact updated Savannah - "massive restagein" >30K files causing delays on many activities incl export FZK-> 3) quicker actions on relatively simple actions, e.g. upgrade SQUID, check these files etc. Good news - some tickets closed (Pisa, Baria, FR T2s, ...) Closing transfer debugging efforts, all details in twiki.
- ALICE (Patricia) - WMS in France back in production this morning & working well. Only very few jobs running over w/e. CREAM nodes at CERN & SLC5 WNs behind working ok too. 3 new CREAM CEs announced today and will be test today/tomorrow. IHEP Russia, SARA and Torino.
- LHCb (Roberto) - WMS outage at CERN - strongly affected activities during w/e. Guess - megapatch - also at RAL. Asked to roll-back to previous version. Developers still not ready to investigate. GridKA: unfortunate error during admin procedure - lost 2K files (cross-check with FZK report below). h/w problem with storage for online DB - once again during w/e. GGUS regarding problem accessing a file on disk; globus xio failures again; Lyon failing to upload files to pnfs; biggest issue - SRM CASTOR at CERN problem preventing normal activity for all users. Thought overload from dummy MC upload -> CERN -> normal users. Opened a GGUS ticket: Gav - still looking into this.

Sites / Services round table:

- GridKA/FZK (Jos) - Finished separating of the ATLAS VO from the the existing dCache SE. During this activity we lost several hundreds of files from LHCb because of operator error. An exact list was send to the local LHCb representatives. GridKa has adapted the procedure for disk migration which will prevent this error from happening again. Awaiting follow-up from LHCb. Angela - LHCb LFC-L SAM tests not running since yesterday evening. Roberto will check.
- RAL (Gareth) - WMS problem Friday evening - became completely unresponsive & m/c reboot. Same as CERN problem? But its post-megapatch. Nick - maybe similar- any details we can look at? Gareth - will get him to do so. Markus - two new open bugs for WMS raised recently. Gareth - we'll look at these first.
- CERN (Andrea) - one of WMS at CERN used by SAM tests shows problems - list match takes alot of time then gives an error - noticed this morning. Markus - did CERN move all WMS to megapatch? Nick - no it was staged as agreed. Steve - not obviously related to megapatch.. Sophie (after the meeting) - all the WMSes at CERN are now running the so-called "mega-patch". And we are having quite a lot of trouble with them... I've asked the EMT to track the following WMS bug?<https://savannah.cern.ch/bugs/?47040>
- DB (Maria) - scheduled maintenance of LHCb online - organised by LHCb sysadmin. Evaluating outstanding issue with Oracle bug traced by CASTOR team but no evidence on DM services.

- CERN (Ricardo) - all actions from post-mortem on CASTOR ATLAS have been done; downtime in GOCDB, CASTOR DBA contact etc. Tomorrow's intervention will not apply patch for CASTOR ATLAS stager - waiting for combined patch. "We maintain tomorrow's slot but only to deploy the patch on the SRM rac. We leave the stager for when we have a merge patch including the bug fix applied this weekend."
- GRIF (Michel - by email after the meeting) - on our WMSes we are running the mega-patch and we also have quite a lot of problems, in particular with WM crashing.

AOB:

## Tuesday:

Attendance: local(Gavin, Alessandro, Jamie, Ewan, Maarten, Sophie, Jean-Philippe, Nick, Andrea, Roberto, Patricia, Julia);remote(Gonzalo, Jeremy, JT, Angela, Gareth).

Experiments round table:

- ATLAS (Alessandro) - yesterday problems TRIUMF-FZK channel - quickly solved by TRIUMF. Long thread about BNL VO box. DQ2. Updated BNL Panda site - marked as SRM V1 - now marked as SRM v2. Planning to integrate BNL VO box monitoring into SLS monitoring of vo boxes (as for other clouds). Scheduled downtimes for PIC & RAL - should finish RAL today, PIC tomorrow. Reprocessing: site validation still ongoing. When completed another round of signoff of monitoring histograms from detector experts. JT - seem to be at least 4 distinct classes of jobs for ATLAS production, all same credentials. 2 concerns: 1) sometimes gets in mode that job completion rate is 8-10 per minute (from ~1 normally). Are these pilots that don't pick up payload? Something else? 2) ATLAS jobs maxing out network bandwidth - upgrade planned for next month. Ale - will check. 15" CPU time jobs - most likely pilot jobs that don't find any payload. Running all jobs with same credentials is not optimal - will check. Working in glexec direction. Jeremy - noticed that VOID card in CIC portal said 100GB for s/w shared area - maybe should be 160GB - 200GB. Brian - getting many jobs failing due to memory at RAL: downtime was an "at risk" for 3D DB so most of transfers for UK still working. Ale - saw outage on GOCDB - to be checked. Mismatch in memory requirements of jobs - should be solved now!
- CMS reports - no attendance due to ongoing CMS week. One item we have is a refreshed request for ELOG for CMS, see mail sent this morning to James Casey. The CMS contact for this is Patricia Bittencourt (cc'ed), I would appreciate if you could kindly please provide her all the needed support in this.
- ALICE (Patricia) - ramping up production again. VO boxes at CERN not able to renew proxies of users - proxies expired although corresponding service running on these m/cs. Will check this afternoon. WMS 215 was performing slowly - jobs in status waiting for a long time. Sophie - maybe this is related to fix applied to all WMS machines. Patricia - also using 214. Sophie - daemon on 215 crashed and was restarted with a buggy script.
- LHCb (Roberto) - Unit test at CNAF: problem understood by site admins - wrong setup of gridftp servers. SRM CASTOR issue reported yesterday - investigation still in progress. WMS: both instances at CERN now working. Sophie - fixed workaround today at 13:00. Ewan - might expect other issues. Have prepared 2 nodes with last release before the megapatch - these can be used if required.

Sites / Services round table:

Tuesday:

- WMS (Maarten) - one known issue with release with simple but annoying workaround. This could be - and was - automated at RAL and now also at CERN. This problem is therefore fixed 'automatically'. Run fix manually on LHCb m/c in debugger. This kept failing on one of the jobs - i.e. even with this workaround in place the service could not recover automatically. Admin intervention required - had to move a job out of the way. Some new bugs opened yesterday. Discussing with developers if fix from that could be back-ported from 3.2 code to 3.1. 3.2 could be many weeks or even months away! Could have more workarounds - cron jobs etc. No definitive statement yet. For some VOs previous WMS was 'better behaved' with workarounds in user space. This could imply 2 versions at CERN - previous version was worse for ALICE! Release apparently all working fine at CNAF - at least for CMS! JT - should not rely on CNAF experience for these services! (maybe a developer is patching the production service!) Nick - still trying to chase down h/w so we can replicate issue and give developers access to debug. Andrea - is this same problem as seen by SAM? A: yes. Maarten - 3.1 branch is now 6 - 12 months old. Sophie - please each VO check that the current system - megapatch plus workaround(s) - is ok so that we don't have to downgrade. JT - we get this an awful lot! Fixes in new branch and can't be back-ported. Never get to stable situation... Jeremy - statement today in a GridPP was that 'WMS is killing Grid adoption'.
- NL-T1 (JT) maybe one related to network max out. BDII drops out every so often. Error code is 80(implementation specific error). Nothing in logs - lsblogmessage failed followed by succeeded! Stumped - ideas? Maarten - which version? One on PPS which many sites have already taken. One had similar error to this. (JT - 4.0.1-0) Maarten - should have -4. "Reducing memory footprint".
- FZK (Angela) - pool manager for dcache went out of memory over night and had to be restarted.
- CERN (Gavin) - DB intervention on CASTOR ATLAS SRM and on CASTOR CMS SRM & stager - both transparent! Sophie - strange permissions on LHCb directories - inherited from top dir. Jean-Philippe & Roberto will check.

AOB:

## Wednesday

Attendance: local(Sophie,Ewan, Antonio, Nick, Jamie, Jean-Philippe, Roberto, Andrea, Julia, Daniele, Miguel, Gavin);remote(Jeremy, Angela, Gareth).

Experiments round table:

- ATLAS (Alessandro) -
  - ◆ Central Catalog problem: dead for the DDM Tracker calls, plus many restarts were done. root cause: upgrade on monday, Birger Koblitz did the rollback.
  - ◆ FZK srm problems (ggus 47198)
  - ◆ Functional Test weekly results: for all Tier1s 100% except FZK-LCG2\_DATATAPE 98% ; TRIUMF-LCG2\_DATADISK 99% ; (and, of course, TAIWAN)
- CMS reports (Daniele) - top 3 site-related comments of the day: 1. one file not going to Castor@CERN tapes since >12 days ! 2. mostly progresses today on T2 issues (maybe a CMS-week effect..) 3. as from yesterday: some sites need to take quicker actions on some relatively simple items (guided upgrade, check list of few files, ...) One broken tape at GridKA caused 4 files to be invalidated - have to be re-transferred. Taking quite some time on actions for a relatively small number of files!
- ALICE -

- LHCb (Roberto) - issues related to default umask on WNs at NIKHEF. Too restrictive permissions on creation of directories in LFC - clients inherit mask from o/s where they are running. Q - why is o/s mask inherited? A: can force a mask. man lfc-python. 2nd issue with umask concerns glxexec test. working directory for pilot jobs too restrictive - pilot can't run jobs once new payload arrives. Workaround in pilot wrapper sets umask less restrictive. For NIKHEF will rerun 100 jobs against NIKHEF CE to understand why FEST activity had problems. Disk server issue this morning at CERN - resolved by Miguel. Miguel - not a diskserver issue, ripple of transparent security patch (Oracle)?? Could not find anything wrong - will investigate abit more. Anomalous # of jobs failing at Lyon T2 being killed by batch system as mem exceeded by job (> VO ID card). Problems accessing files at NIKHEF due to 2nd gplazma server started by mistake. Now ok. GGUS ticket against CERN for pilot role - mapped to different account. Debugging session with Gavin on srm timeout problem reported the day before - srm daemon didn't pick up req. from srm db - all threads busy or stuck(?) Gavin opened service request to Shaun to investigate + enhancement to logs. Sophie - LFC stuff - don't you create user directories from here? A: yes, part of new user creation - new user in VOMS -> run script. How does this figure with umask at NIKHEF??

Sites / Services round table:

- ASGC- ASGC T1 and Taiwan Federated T2 services will be collocated at IDC from Mar. 19. All the T1 and T2 services planned to be up and running before Mar. 23. We hope to make 2,500 cores and 1.3 Petabyte disk space available next week. Tape Library will be available about one more week later with clean tapes added gradually. During the transition, all ASGC T1 and T2 services will be shutdown and restart at the last day (Sunday, Mar. 22).
- RAL (Gareth) - ATLAS thought we were in an outage but we weren't... Do look at ATLAS Grid downtime - google calendar - seems to show outages that are nothing to do with ATLAS! Ale - we have script that parses RSS feed from CIC portal - should check what was published in RSS feed.
- Release update (Antonio) - glite 3.1 pps update 45 to pps is progressing - passed deployment tests. Contains 2 FTS patches for fixing FTS job submission that sometimes ends up with invalidated delegated proxy. Can't yet forecast release to production as others in pipeline: 1) glite 3.2 update 01 - 23 Mar, SL5 WN 2) glite 3.1 update 42 - new VDT version - affect all service nodes. Around 30 Mar. 3) 6 Apr CREAM release - ICE + CREAM - version "blessed" by pilot. Contains 5 patches. CREAM & WMS - much better than one currently in production. Nick - will suggest through Markus to MB on milestone for CREAM CE ready to start using this release. Antonio - this would be first usable version of CREAM wrt submission through WMS (and not just direct submission). Andrea - bug in Savannah on WMS submission in pps? Bug in Savannah by developers...

AOB: (MariaDZ) Latest web page with tomorrow's CERN Disruptive Network Intervention March 19th 2009 details: 090319-Network-Intervention [HERE](#).

## Thursday

Attendance: local(Diana, Jamie, Nick, MariaDZ, Patricia, Maria, Gavin, Alessandro, Roberto, Ignacio, Ricardo);remote(Gonzalo, Angela, Gareth).

Experiments round table:

- ATLAS (Alessandro) - checking behaviour of services (after network intervention) - seems all up and running. Small problems around 10:00 from site services machines - increasing number of callbacks to dashboard. Error - restarted consumer and solved. Asked for possibility for ATLAS experts on-call to be able to do such operations - Savannah bug request. ~11:00 many critical errors on site service m/c - not working central catalog. Interaction with Oracle DB not stable? Contacted Ulrich - understood how to restart central catalog and since 13:00 ok. Stopped functional test load generator

this morning and will restart this afternoon.

- CMS reports - (apologies from Daniele - clash with CMS meetings)
- ALICE (Patricia) - since yesterday production stopped at many sites due to WMS problems. 214-5 jobs in waiting for ever, 221 in draining mode. Nick - symptoms as before megapatch? Yes, but monitoring doesn't show any overload. 15' ago myproxy server not responding - Gav: myproxy stopped for about 1h and should now be back. Will send e-mail to ALICE-taskforce list to explain.
- LHCb (Roberto) - network intervention: LHCb only rely on one non-distributed service: LFC\_write. Reminder not sent to (LHCb) users so a few unaware... From shifter: status update of all T1s: Lyon - serious dcache config problem that prevents first pass recons as locality reported by SRM is wrong. Lionel opening dcache service request. NIKHEF: some problems accessing files - only for prod jobs and not users(?) Prod files sitting in different (RAW) space token. Maybe something to follow up? Other T1s ok. SAM jobs and user jobs running smoothly. Wrong mapping of user & pilot roles at CERN - seems to be true also at all T1s! Will check... Next week: on-going analysis activity; very low level of dummy MC production - 2K jobs per day.

Sites / Services round table:

- ASGC (Jason) - minor update:
  - ◆ the delivery of the tape system would be difficult to be online in one week due to the cleaning of existing tape system is not well fit with schedule proposed. besides, all tape drives are not able to cleanup due to the complex interior components of the LTO3/4 drives. the MES rocurment expect to finish mid of next week, while the actual deliver term will extend another 45 days. though vendor promise to speed up the internal bidding as well as shipment but still might delay another two weeks or more. local IBM promise to loan for another LTO4 drive but we might have limited b/w if two VOs request the migration/stagin around the same time. we're still negotiating with local vendor see if having chance expending the drives to two or three more.
  - ◆ migration of the data from existing cartridges will be another concern while we need to confirm other technical details, procedures before taking any action.
  - ◆ merging/splitting of TS3500 tape system will delay another 3-4 days, depend on the labor order from the MES case.
  - ◆ hope for your understanding, and we hope to relocate also the tape system into collocation data center area while we need better coordination due to the limited floor plane serving the tape library. Alessandro - is there any foreseen time when they will be ready? How could ATLAS handle situation with a T1 down so long? Trying to move Australian T2 to a different cloud. Maybe use LFC and FTS in TRIUMF?
- RAL (Gareth) - planning to turn on "scrubbing" of disks on CASTOR on disk servers (RAID systems on these). (Ignacio - scrubbing foreseen for castor 2.1.8)
- CERN (Gavin) - status of CERN services: data services (CASTOR&FTS) went down ~06:00 at start of intervention. FTS channels set inactive and will be reopened around 16:00 when CASTOR back. WMS services drained around 05:00. Batch also drained. WMS back at 10:00. Batch now accepting jobs but not processing them! Queuing but waiting for castor to be back online - all should be back around 16:00. Myproxy and BDII - did some tests & reboots. Top BDII down for about 75' from 10:45 - 12:00. Myproxy failover tests - left off from 12:30-13:30. (To find dependencies). Found at least one dependencies on FTS - review at workshop!
- DB (Maria) - stopped 05:45 validation and production clusters on GPN; (ATLR,CMSR,COMPR,LCGR,LHCBR,PDBR,ATLDSC, LHCBDSC, INTR, INT2R, INT6R, INT8R, INT9R, INT11R, INTR12R, D3R) brought up at 08:00. Profited of downtime to apply some

parameter changes to fix Oracle clusterware timeouts (which caused reboots). ATLAS DB situation of ASGC: contacted by Jason that ATLAS DB back up. Since a few days listener problem - no answer. If it cannot be restarted will have to remove from replication setup - queue memory 75% full.

AOB: (MariaDZ) The Action Plan for the GOCdb-to-OIM move out of GOCdb into OIM <http://indico.cern.ch/getFile.py/access?resId=0&materialId=minutes&confId=54492> was sent to all yesterday. Please be patient during the migration period (ending this May) for the multiple sitemames one sees today behind "Affected SITE" on <https://gus.fzk.de/pages/ticket.php> Any misrouted tickets, use the escalation button or bring them to this meeting.

## Friday

Attendance: local(Nick, Roberto, Jean-Philippe, Alessandro);remote(Jeremy, Gareth, Brian).

Experiments round table:

- ATLAS (Alessandro) - service up and running ok after network intervention. Functional test restarted late afternoon. Still problems with FTS delegation at FZK (proxy corruption). ATLAS proposes to use alarm tickets in this case to push sites to install FTS pre-release. Reprocessing has not started yet. May be this weekend but more probably on Monday (James Catmore responsible for it). Nick asked if any problems seen with WMS at CERN. Answer: not after the workaround has been applied. In correct monitoring at RAL? RAL sees 4 Gbits/sec out of MCDISK: not seen by Alessandro.
- CMS reports -
- ALICE -
- LHCb - fix problem of mapping to pool accounts at CERN: use pool accounts for requests without WOMS role, fixed accounts for requests with VOMS role. 13000 files lost at NIKHEF due to a pool node. Entries have now been removed from PNFS. Relatively low activity: 2000 job/day grid-wide.

Sites / Services round table:

- RAL (Gareth): top level BDII problem at 12:00. Root cause not understood yet. Has any other sit seen similar problems? At RAL they stop the BDII service, cleaned the cache and restarted the service. Nikhef has also seen a BDII problem a few days ago. The corresponding ticket number is 47177. Will check if similar problem.

AOB:

-- JamieShiers - 13 Mar 2009

---

This topic: LCG > WLCGDailyMeetingsWeek090316

Topic revision: r14 - 2009-03-20 - JeanPhilippeBaud



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