

Table of Contents

Week of 090727.....	1
WLCG Service Incidents, Interventions and Availability.....	1
GGUS section.....	1
Daily WLCG Operations Call details.....	1
Monday:.....	1
Tuesday:.....	2
Wednesday.....	3
Thursday.....	5
Friday.....	6

Week of 090727

WLCG Service Incidents, Interventions and Availability

VO Summaries of Site Availability				SIRs & Broadcasts	
ALICE	ATLAS	CMS	LHCb	WLCG Service Incident Reports	Broadcast archive

GGUS section

- GGUS Escalation reports every Monday (used for WLCG Service Report to MB)
 - ◆ Full search: https://gus.fzk.de/ws/ticket_search.php and select "VO" - this will give all tickets, including team & alarm
- LHCOPN Tickets in GGUS: see https://gus.fzk.de/pages/all_lhcopn.php and change your selection criteria. Future actions are also listed.
 - ◆ If network group participation is necessary, please invite them in time.
- Procedure to become a LHC Experiment VO TEAM member
- Other recent GGUS FAQs

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			
CERN IT status board	M/W PPSCoordinationWorkLog	WLCG Baseline Versions	Weekly joint operations meeting minutes

Additional Material:

STEP09	ATLAS	ATLAS logbook	CMS WLCG Blogs
------------------------	-----------------------	-------------------------------	--------------------------------

Monday:

Attendance: local(Julia, Jean-Philippe, Ricardo, Eva, Jamie, Harry, Olof, David, Andrew, Alessandro, Gang, Simone);remote(Daniele, Michael, Andrea, John, Angela).

Experiments round table:

- ATLAS (Ale) - not many issues: observed a problem vs DESY - due to misconfig from central OPS. Monitoring showed it and fixed Sat night. Issues with 2 T2s: 1 RO (FR cloud) and 1 RU (NL). Observed late Friday pm security incident - already discussed with security team. This morning ATLAS elog not accessible. James Casey answered - only a few hours downtime - no contacts except one single person: should be improved.
- CMS reports (Daniele) - full details in report! (Follow link) ASGC(Gang): for data migration due to streaming problem(?), for link connection quality waiting for reply from Russian T2 admin.
- ALICE - RAS

- LHCb reports - (not updated - no report)

Sites / Services round table:

- FZK(Angela) - informed in morning about security intervention tomorrow morning. Short network interruptions between 7 and 8 CEST. Only informed late - sorry for the late announcement. Declared "at risk" at this time.
- BNL (Michael) - ATLAS farm will be upgraded tomorrow to SL5. Other services will be performed so expect to be down for the entire day. Already announced in GOODB.

AOB:

Tuesday:

Attendance: local(Ricardo, Eva, Jamie, Olof, Simone, Jean-Philippe, Andrea, Gavin, Julia, Gang, Harry, Roberto);remote(Tiju Idiculla (RAL), Xavier Mol (FZK), Ronald Starink (NIKHEF / NL-T1), Michael, Daniele Bonacorsi [CMS], Jeremy).

Experiments round table:

- ATLAS - 1st: problem at CERN with acrontab - users A-L. DDM users in this range. No interruption of service for DM & WM but monitoring affecting (SLS relies on acron). VO Boxes flagged as "downgraded(?)" - 30% availability. Now cured. Functional tests - which rely on batch system at CERN - also affected (same reason). acron back in service mid-late morning. 2nd: dashboards - reminders / alarms from monitoring system about ATLR DB & dashboard app - locked sessions. Got a variety of alarms in the morning - can someone from dashboard team check? Julia - nothing seen on support list. Simone - will forward one. 3rd: ASGC status: reprocessing tested. Seems to be no more workload management problem. Jobs can run at T1 - peak at >1K jobs. Remaining issues: test stopped as CASTOR showing XIO errors. Failure rate too high so test stopped. Ran 4K reprocessing jobs -quite good! Main worry: ATLAS cannot run when CMS is running & vice versa. Access to tape system. FIFO access to tape requests & not enough tape drives. Makes little sense to run ATLAS only exercise - need ATLAS+CMS. Files lost in TRIUMF - thousands! MC files - can be regenerated. Very little interest. Believed from SRM 1-2 migration. Cleaned from catalogs. Yesterday reported unavailability of 2 files (CASTOR) at CERN - problematic disk server- understood. Yesterday DDM bug - happens when unavailability of Oracle b/e etc. Client side: client not transaction safe. Transaction not rolled back - caused troubles 20days ago - problem will be fixed. SARA is in unsched downtime since this morning. What is the problem? Last point: FTS - noticed that transfers CERN-BNL failing as channel not defined in FTS at CERN. Looks like BNL stopped publishing sitename under old name. Michael- part of site name consolidation. To surprise of experts here site name change affecting transfers such that they fail. Name change reversed but will go ahead to work towards consolidation. BNL in sched down for entire day anyway. Should not worry about these failures at this point in time. Gavin - understand go back to BNL_LCG2. Michael - yes, then BNL_ATLAS1 in future. Gavin - Will affect CERN and all T1 FTS servers. Can generate procedure for this.
- CMS reports - 1st: on dashboard had issue on DESY CE not properly included. Fixed. ASGC: still backlog of migration. ~2.5TB of data to be migrated to tape. Files ready for migration - tape drives free but migration not proceeding. Expect some news later... Some issues at PIC: CMSSW had some R/O f/s errors but these are solved. T2s: all tickets open apart from one at Caltech. >3/4 related to intense T2-T2 commissioning activity. Julia: need to upgrade schema for dashboard job monitoring - request to track output SE and some parameters of staging out. Probably need to coordinate this upgrade. Daniele - stable production from now until early September! Please suggest a duration for this for facilities hypernews. Some midday slot probably fine. Just propose a date... Andrea - acron

Monday:

affected also CMS: PhEEx, DBS and FroNTier.

- ALICE -
- LHCb reports - 2 weeks ago reprocessing activity run after STEP'09 to test staging capacity: post-mortem made available, see LHCb report page. As for ATLAS & CMS. acron problem (SLS sensor & SAM test submission). A lot of MC production - up to 10 MC prods from physics groups. 15K concurrent jobs. Sustained for several days. 2 different problems, 1st at CERN reported yesterday with LHCb data pool - full - more disk capacity requested. Bernd agreed to provide 100TB. Got 5TB temporarily... CNAF - sqlite DB file in shared area again? - looks like fix to copy file locally to WN before job doesn't work for all s/w versions, site banned from production. Will chase problem up via GGUS ticket.

Sites / Services round table:

- FZK - this morning's services will be at risk for short network interruption performed without problems. Did some security upgrades on dcache. Also no problems, production not interrupted.
- NL-T1 - SARA started dcache upgrade this morning. Last info: SRM nodes wouldn't start. Will continue to work on this. If cannot be solved before tomorrow will be downgraded again. Version # 1.9.3-3.
- GridPP - redoing Hammercloud tests that ran last week with some problems.
- ASGC - globus XIO error as ATLAS pool only has 1 server. Has reached its limit. Waiting for another 2 servers expected mid August. Tape drives: # jobs sent based on # CPUs. If other resources cannot be increased as fast this will be a problem - big increase in CPU resources recently. Simone - should rerun tests after mid August - TBC with Jason.
- DB - LHCb online DB down since ~21:00 last night. Someone pressed a red emergency button - everything stopped! Trying to recover DB from backup. SLC5 upgrade postponed - problem understood, testing the solution before going ahead with migrations.

AOB:

Wednesday

Attendance: local(Julia, Gavin, Eva, Antonio, Harry, Jamie, Olof, Ale, Edoardo, Jean-Philippe, Gang, Simone, Ricardo);remote(Ronald, Michael, Andrea, Xavier, Tiju, Daniele).

Experiments round table:

- ATLAS - email from Simone:

1) BNL is back in production after yesterday dCache upgrade. All configuration changes associated with the site name consolidation had to be reverted, i.e. the US ATLAS Tier-1 center is back to BNL-LCG2 (from Michael). Therefore FTSEs can transfer files from/to BNL, but it remains an open question of how to have a smooth transition once the name is changed.

More details as to the results of the investigation (summarized by John Hover (BNL)) became available. The issue was ultimately traced back to the way a particular GIP probe is configured within CEMon. Despite assurances that the change could be made solely in OIM (OSG's GOCDB), there was an inconsistency that made glite-query-sd fail.

Tuesday:

The wider issue is that OSG is moving to a model where "resource groups" correspond to EGEE "sites", whereas previously OSG "resources" roughly corresponded to sites. The OSG tools have not all made this shift consistently, and no one seems to fully understand all the ramifications.

Deeper technical details and troubleshooting discussion can be found at this Twiki page which is being used in place of an email thread:

<https://twiki.grid.iu.edu/bin/view/Operations/AtlasBdiiIssues>

Gavin provided scripts and procedure to change a site name in FTS

<https://twiki.cern.ch/twiki/bin/view/LCG/FtsProcedureSiteNameChange>

It is very good news that a clear recipe and scripts now exist to quickly switch sitenames for FTS. That should make our eventual consolidation much easier.

It must be understood how to do this transparently.

2) SARA this morning was still having problems with the dCache upgrade. The information from Hurng might be interesting for other sites:

""

After the upgrade of dCache, SARA was suffered by a problem that the new version of dCache writes out additional information to the billing file, which then cause the partition of the main dCache node full (instead of 40 MB/day .. it grows up to 5 GB after few hours).

SARA people is working on it and contacting developers for a good solution. Another downtime was claimed.

""

- CMS reports - CMS - progress on T1 link commissioning ASGC-RU site. DPM layout changed. Expect commissioning to progress soon. No other progress on open tickets, e.g. ASGC migration, PIC load test T2 transfer. T2s: quite a lot of progress in T2-T2 link commissioning. High level of activity in US and UK. GRIF/Roma tickets closed - good progress in this activity!
- ALICE -
- LHCb reports - There are currently running 16K jobs concurrently for 6 different physics MC production. In this picture a snapshot of the last 24 hours running jobs. T0 issues: Intermittent time outs retrieving tURL from SRM at CERN (verified for the RDST space, most probably general issue Observed LFC access times out this morning. SLS does not seem to indicate problems. T1 issues: issue with SQLite DB file in the shared area: provided to CNAF people all suggestions to fix this problem as done at GridKA. Waiting for them.

Sites / Services round table:

- SARA (Ronald) - did upgrade due to vulnerability. Try not to make habit of it! This morning and last night had problem with partition filling up with logging info - about orphaned files (on disk but not in namespace). Lost but taking up space. Have to be cleaned up. This new version of dCache complains at quite a high rate. Separate partition for logging but this was written in billing file! May affect sites upgrading from dCache 1.9.1 or older to 1.9.2 or newer! (If site also has a lot of orphaned files).

- FZK - shortly after midnight dcache srm space manager for all VOs apart from ATLAS died for out of memory. Only restarted at working time. Monitoring confused by errors - not always critical, sometimes warning or unknown hence on-call not activated.
- BNL - as reported earlier site name consolidation failed and is under investigation with OSG colleagues. OS upgrade to 64bit SL5, NFS server (split between analysis and production now). Upgrade of AFS server - all completed within announced downtime.
- ASGC - concerning CMS savannah ticket now understood - tape pool assigned is full. Added 80TB of space and migration then started again. About 400 files migrated in last few hours. (Thanks Gang from Daniele!)
- DB - LHCb online DB back in prod since yesterday evening. During last night ASGC ATLAS DB down until this morning - some problems with storage and listener. Waiting for Jason to restart propagation CERN-Taiwan
- SRM b/e DBs (Gav) - patch scheduled tomorrow at 15:00 CEST - should be transparent!
- Release update (Antonio): BDII V5 released last week. Maybe of interest for services on SL5. Some fixes for issue found at one of production sites - affects BDII on SL5 - some sites may disappear due to ongoing change of GLUE schema. If you run SL4 BDII only affect is a corruption of LDAP tree. SL5 top-level BDII should run this fix released glite 3.2 update 04. New version of LFC & DPM on SL5, new gfal libs, myproxy client on WN. For future preparing update glite 3.1 update 52 on SL4. Update to CREAM CE. Will start to be published correctly in IS. New host certs for VOMS server. Current certs expire end August - voms.cern.ch certs to be replaced and will be in next release.
- Network - upgrade DNS at CERN due to vulnerability. No impact except for load balanced alias 07:00 - 07:30.

AOB:

Thursday

Attendance: local(Ricardo, Jamie, Julia, Eva, Jean-Philippe, Roberto, Harry, Gang, Simone, Alessandro, Miguel);remote(Xavier, John, Michael, Daniele, Jeremy).

Experiments round table:

- ATLAS - nothing!
- CMS reports - update on ASGC ticket: migration backlog being digested. Just 7 files waiting now.. Problem seems to be now solved. Some progress especially in T2-T2 link commissioning - details in twiki.
- ALICE -
- LHCb reports - Previous days running at full regime productions have now been stopped because the failover is also getting full at all T1's due to the lhcbdata space at CERN full. CERN disk capacity is indeed today the main problem (also accordingly SLS) and became a show stopper. LHCb would like to hear from FIO which the plans to get these extra 100TB agreed on Monday. Ricardo: machines are being prepared, still have to be installed and configured but should start to become available next week. Roberto: CNAF working with LHCb & Roberto on SQLite issue.

Sites / Services round table:

Thursday

- FZK - same problem as yesterday but this time on-call called as monitoring improved. SRM service restarted... Local monitoring still failing as cert outdated. Normal SAM & OPS tests succeeded - production not affected.
- BNL (John R. Hover [jhover@bnl.gov])

Hello all,

As Michael said, all changes were rolled back so nothing needed to change on the EGEE side.

The issue was ultimately traced back to the way a particular GIP probe is configured within CEMon. Despite assurances that the change could be made solely in OIM (OSG's GOCDB), there was an inconsistency that made glite-query-sd fail.

The wider issue is that OSG is moving to a model where "resource groups" correspond to EGEE "sites", whereas previously OSG "resources" roughly corresponded to sites. The OSG tools have not all made this shift consistently, and no one seems to fully understand all the ramifications.

Deeper technical details and troubleshooting discussion can be found at this Tiwiki page which is being used in place of an email thread:
<https://twiki.grid.iu.edu/bin/view/Operations/AtlasBdiiIssues>

It is very good news that a clear recipe and scripts now exist to quickly switch sitenames for FTS. That should make our eventual consolidation much easier.

Cheers,

--john

- LHCb - also ticket for SRM problem at CERN due to disk space full. Fixed now.

AOB:

Friday

Attendance: local(Luca, Harry, Roberto, Alessandro, Jean-Philippe, Ricardo, Gavin, Gang);remote(Xavier/FZK, John/RAL).

Experiments round table:

- ATLAS - 1) Problems staging data in at RAL with some files waiting since 26 July. Ticket raised and is being worked on. 2) Planning to add a third machine into the central catalogue pool next Monday. Should be transparent. 3) The LCG tags interface does not work on SL5 worker nodes so ATLAS cannot update their WN software on SL5. A new version 0.4 that does work will be released with gLite 3.2 so ATLAS will meantime package this with their software.
- CMS reports - Apologies for absence today.
- ALICE -
- LHCb reports - 1) Monte Carlo production has been put on hold due to the LHCBDATA space token being full at CERN. An increase is expected next week - an advance warning would be appreciated. 2) IN2P3 is publishing CE names in a way that prevents LHCb agents from working. 3) Two UK Tier 2 sites (Lancaster U. and Queen Mary London U.) have a high job failure rate.

Sites / Services round table:

- * FZK: SRM out of memory problems and associated 'confused' monitoring both now fixed.
- * RAL: The lcg monbox disk replacement yesterday during a scheduled outage subsequently failed leading to an additional unscheduled outage.
- * CERN Databases: The LHCb Online database will be down next Monday and Tuesday for a scheduled maintenance at P8: "upgrade of the core-router in the SX8 network and the electrical installation".

AOB:

This topic: LCG > WLCGDailyMeetingsWeek090727

Topic revision: r10 - 2009-07-31 - HarryRenshall



Copyright &© 2008-2021 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](#)