

Table of Contents

Week of 100125.....	1
LHC Operations.....	1
WLCG Service Incidents, Interventions and Availability.....	1
GGUS section.....	1
Daily WLCG Operations Call details.....	1
Monday:.....	1
Tuesday:.....	3
Wednesday.....	6
Thursday.....	7
Friday.....	9

Week of 100125

LHC Operations

- [Sharepoint site](#) - [Cooldown Status](#) - [News](#)

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.
- OSG items selected from the GGUS escalation reports
- 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](#)
3. The scod rota for the next few weeks is at [ScodRota](#)

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(Jaroslava, Jamie, Harry, Ignacio, Julia, Timur, Nicolo, Eva, Patricia, Roberto, Miguel, Steve, MariaDZ, Gavin);remote((Jon Bakken(FNAL), Kyle (OSG GOC), Rolf (IN2P3), Michael (BNL), Gonzalo Merino (PIC Tier1), Gang Qin, Angela Poschlad (KIT), Rob Quick, Gareth (RAL)).

Experiments round table:

- ATLAS reports - Weekend issues: srm-atlas.cern.ch SRM v2.2 endpoint timeouts, was red on SLS monitor, high request load on _SCRATCHDISK, alarm GGUS:54949; No transfers between CERN and NIKHEF: CERN FTS restarted on 23 Jan, further details promised during Monday,

GGUS:54935 [↗](#); Slow transfers between Triumpf and ASGC, got better during 23 Jan; Observed _MCTAPEs short of space (BNL_OSG2, TRIUMF_LCG2, PIC), not sure whether garbage collector will be fast enough under load or do we need larger _MCTAPE buffers. Steve - ATLAS NIKHEF downtime - see CERN report below.

- CMS reports - Highlights:
 - ◆ T0
 1. T0 operators lost interactive access on cmst0 worker nodes (needed for debugging troublesome workflows). T0 ops reports successful login, requesting additional access from vocms68 and vocms69. GGUS #54848
 2. Saturated T0 with 2700 jobs running on Friday, some RFIO open errors: 'Job timed out while waiting to be scheduled' - maybe correlated to high number of queued transfers on TOEXPORT.
 3. New SLC5 VOBOXes vocms02, vocms03 available for PhEDEx, requested registration in myproxy Remedy #CT656082. PhEDEx Debug instance migrated to SLC5, checking for issues before migrating Prod instance.
 - ◆ FNAL
 1. New ReReco pass ran at FNAL.
 2. One very large (120 GB) file from a test replay timing out in transfers T0-->FNAL.
 - ◆ PIC
 1. JobRobot errors during weekend: MARADONA error.
 2. File waiting for a long time to be exported to T2s.
 - ◆ ASGC
 1. Following up on deployment of SLC5 software releases
 - ◆ CMS T2 sites issues: Highlights (selected ones):
 - ◇ Various MC production workflows running in KIT, IN2P3, RAL, CNAF, PIC, FNAL T2 regions - T2_FR_GRIF_LLX excluded for scheduled downtime.
 - ◇ Open tickets for T2_EE_Estonia, T2_BR_SPRACE
 - ◇ T2_UK_SGrid_Bristol migrated production storage element to StoRM.
 - ◇ T2_BE_UCL - SLC5 nodes configured. T2_UK_London_IC - requesting installation of dependency meta-RPM on SLC5.
- ALICE reports [↗](#) -
 - ◆ T0 site
 - o On Friday ops meeting a kernel upgrde was announced (performed on Thursday, nodes were drained and rebooted on Friday). The operation was announced as almost done at the T0, however during the whole weekend and also today, SLC5 batch systems are announcing 99 total CPUs. Issues to be aksed to the expets today during the ops. meeting o Access to the ALICE VOBOXES reported last week, still having problems with voalice14. Access to this machine is immediately closed while it is using the same CBD templates as voalice11, 12 and 13. Remedy ticket opened yesterday: CT656294. This is concerning the vobox submitting to CREAM-CE, the issue should be considered of high priority
 - ◆ T1 sites
 - o CNAF T1: On regard with the CREAM-CE, the corresponding queue was reporting zero CPUs at the site yesterday afternoon. Reported to the site admin and corrected this morning. Regarding the submission to the LCG-CE via gLite-WMS a missconfiguration in LDAP was avoiding the proper behavior of the Packman service. Solved this morning. Finally the site is finishing the required configuration of ALICE to test glxexec at the site
 - ◆ T2 sites
 - o CREAM-CE systems announced in two french sites entering production in a short time. Systems announced yesterday for:GRIF IPNO and IRFU o Annotation in Savannah telling CREAM for SGE is being certified: <https://savannah.cern.ch/task/?9126#comment13> [↗](#). This will allow Trujillo T2 to provide this service for Alice as soon as the system is ready in production. o LPSC (Grenoble) vobox not reachable yesterday. Issue solved this morning, services restarted and back in production o Poznan VOBOX suffering of some hardware problems announced this morning. Waiting for the confirmation of the site admin to put the system back in production o Madrid WMS system was showing a bad performance yesterday afternoon (jobs were not submitted to the

queue although the site was empty). Reported to the site admin. WMS services restarted.

MariaDZ - why don't you use GGUS directly to site? A: site initiated contact but we can do this via GGUS.

- LHCb reports - Only bb and cc inclusive MC09 stripping in the system now (very few jobs in the system in total at T1's). Launched a MC simulation but some application problem found. T1 sites issues: RAL : Downtime; IN2P3: Downtime.; PIC: an issue with one file for some user analysis jobs Under investigation by our local contact person. CNAF - migration from CASTOR to TSM now completed and registered in catalog.

Sites / Services round table:

- FNAL - 1) one of BDII servers out of date - not providing correct info. Reported last week - ongoing problem. Ticket was closed & problem fixed but reappeared and reopened. 54803 GGUS. Be good if all those BDII's would report consistent info - if a site gets the wrong one it gets wrong info. 2) OSG reported accounting correction - was refused. OSG GOC 7986.
- IN2P3 -confirm still in downtime will continue until tomorrow. At beginning of downtime a lot of ALICE jobs waiting. Operator in charge made a mistake - about 1/3 ended prematurely. Sorry. Hope others will go through ok.
- OSG - reviewed escalation report. 7 GGUS tickets - 4 closed and some for a week or so. When is this report run? (Monday) Maria - on index page there are many different reports. Tickets assigned to OSG - see documents behind ROC - OSG appears as a ROC. No open issues this week!
- BNL - minor issue: failure of storage server yesterday. For 3 hours job trying to access input files on that server were unable to get them.
- PIC - had some hiccup with SRM server this morning. Causes still not clear - suffering from overload and some transfers timeout 09:00 - 11:30 this morning.
- KIT - ntr
- ASGC - minor update - performance degradation - continue installing 800TB. (of disk?)
- anon -
- RAL - 2 things: 1) currently "at risk": for UPS work - hope of fixing noise on current. (Today - ongoing), 2) also draining batch for intervention Wed/Thu migrating Oracle DBs back to original disk arrays.
- CERN FTS: The CERN-NIKHEF T0 FTS channel was down on Saturday 23rd from 04:09 till 23:53. The CERN operators did intervene earlier but their restart attempts failed. The service manager restarted agent to correct at 23:53. Log analysis gives no clues. Was also reported by ATLAS. GGUS:54935 [↗](#).
- CERN - ATLAS SRM incident yesterday ticket # 54949. Doing a post-mortem. User accessing ATLAS scratch at a high rate - queuing building up. Access from SRM timing out. This can take much more than 3' of FTS timeout. If using SRM to access pool user access should be limited to avoid queuing. xrootd access - unscheduled - would not have shown this behaviour. Would have worked better. (ATLAS - in contact with user).

AOB:

Tuesday:

Attendance: local(Jamie, Gavin, Steve, Jaroslava, Harry, Eva, Lola, Nicolo, Roberto, MariaDZ);remote(Jon Bakken, Gonzalo, Angela, Gang, Ronald, Jeremy, Michael, Rolf, Pepe, Jason Rob).

Experiments round table:

- ATLAS reports - One issue: problems transfers to IT T1. Need to restart SRM daemon and since lunch transfer work well.

Tuesday:

- CMS reports -
 - ◆ T0
 1. T0 operators lost interactive access on cmst0 worker nodes (needed for debugging troublesome workflows). T0 ops reports successful login, requesting additional access from vocms68 and vocms69. GGUS #54848
 2. New SLC5 VOBOXes vocms02, vocms03 available for PhEDEx, requested registration in myproxy Remedy #CT656082. PhEDEx Debug instance migrated to SLC5, checking for issues before migrating Prod instance.
 - ◆ FNAL
 1. Some merge jobs in ReReco pass at FNAL were stuck, dCache intervention by FNAL admins.
 - ◆ ASGC
 1. Following up on deployment of SLC5 software releases
- ◆ T2 sites issues: Highlights (selected ones):
 - ◇ T2_UK_SGrid_RALPP, T2_EE_Estonia, T2_IN_TIFR - ongoing errors in SAM tests.
 - ◇ T2_ES_IFCA - problem in job reporting to Dashboard from production jobs
 - ◇ T2_UK_London_IC SLC5 software deployed, now following up on T2_FI_HIP.
 - ◇ T2_IN_TIFR - network issue solved by GEANT, now good upload rates.
 - ◇ T2_PK_NCP - CERN-NCP channel demonstrated good stability, requested increase in number of files to saturate site bandwidth.
- CMS Weekly-scope Operations plan

[Data Ops]

- ◆ Tier-0: taking data in MWGR (mid-week global run) Wednesday/Thursday and otherwise doing replays and transfer tests to T1. Tier-1: waiting for new rereconstruction requests of 2009 data at custodial sites possibly coming this week, backfill at IN2P3 and CNAF as preparation, more testing everywhere especially FNAL. Tier-2: ongoing MC production

[Facilities Ops]

- ◆ Managing VoBox requests for central servers at CERN. Migration of central services to SL5.
 - ◆ Several Tier-2 sites still have only SL4 worker nodes and several more have SL5 WN's, but no SL5 builds due to various problems. Tickets are opened for them and progressing.
 - ◆ Follow-up the SL5 WNs migration and tape recycling/repack in ASGC Tier-1, to bring the site back to operations and be ready for 2010 run.

Note: Beam Commissioning 09 computing post-mortem session scheduled today at 4.30pm. Discuss lessons learned during 2009 data taking.

- ALICE reports [↗](#) - GENERAL INFORMATION: New MC production cycle, normal running with more than 13K concurrent jobs
 - ◆ T0 site o Issue reported yesterday on regard with the access to voalice14 still following it together with Steve
 - ◆ T1 sites o CNAF: Site admin reported yesterday a large amount of Alice agents being submitted to the site (over the limit defined by Alice to avoid any queue overload). It was reported to the site admin, that the information provided by voview regarding Alice queues was not reflecting the real status of the queue. Site admin is following the issue.
 - ◆ T2 sites o Issues reported yesterday regarding Madrid and Grenoble solved. o Still waiting for a fix on the Poznan VOBOX (system seems to be continuously overloaded) o The regional expert in Italy reports: + Cagliari: The new VOBOX is suffering of some instabilities in the

proxy renewal mechanism. Being followed by the site admin + Catania: The CE at the site is showing a bad performance. Issue being followed by the site admin

- LHCb reports - No scheduled activities running in the system now.
 - ◆ T0 sites issues: o CASTOR upgrade this morning. o VOBOXes:forced reboot this afternoon for kernel upgrade o volhcb15: (LHCb log SE service). Delivered, the machine seems to have not the right partition (just one for "/" and usually ones for OS despite has been explicitly asked a different partitioning)
 - ◆ T1 sites issues o RAL : Downtime o IN2p3: Downtime. o PIC: SE was banned due to many users affected by a dCache pool which had some network problem. The content of the problematic pool has been migrated to a new one. o CNAF:CREAM CE problems if using sgm account.
 - ◆ T2 sites issues: o Shared area issues both at UKI-SOUTHGRID-RALPP and AUVERGRID

Sites / Services round table:

- FNAL - As Nicolo said one problem with pool which had system disk replaced. Incomplete install of dCache s/w - repaired last night and now ok. Rolling upgrade of dCache pools to newer version to fix one of bugs in 1.9.5-11 series.
- PIC - ntr
- KIT - Need maintenance on part of one tape library - plan at risk Monday 08:30 - 12:30. During this time 1/3 of old data will not be accessible for reading. ATLAS has a downtime at this time and local CMS say 'OK!'
- NL-T1 - At NIKHEF have completed kernel upgrade on WNs, UI and VOBOXes. SARA has experienced CREAM CE crash - increased FTS timeout to one of Russian T2s.
- BNL - Issue with firmware on core switch - Force10 - communicated to Force10. They have provided a fix which should arrive today. Will fix asap. Requires reboot of switch. T1 activities will stall for 5-8'. May see a couple of failed transfers but no disturbance to production jobs.
- ASGC - ntr
- OSG - a couple tickets open on BNL. 15089, 15121. Both opened last week(end?).
- GridPP - Observation: increasing number of RAID-related server problems at T2s. Incident at Glasgow - more user analysis. These groups must make copies / backups of files. Reminder!.
- IN2P3 - still in downtime. Small extension due to unannounced update to WNs. Else all on schedule. Storage ok, Oracle security patch underway. Batch controller already there but no batch for moment...
- RAL - We wish to remind everyone that we have an intervention planned for the LFC tomorrow and Castor tomorrow and Thursday. We are also in the process of draining the batch system.
- CERN - ALICE mentioned capacity decrease following kernel update. Corrected this morning. ATLAS also noticed it - ran out of SLC5 resources completed. Now ok.
- CERN DB - ALICE Has upgraded firmware of 3 disk arrays. Still running on DB at CC. Will schedule switchover for end of this week - still TBC. Start deploying latest Oracle security in production this week: Wed - ATLAS; Thu - LHCb. Rest next week. Rolling intervention.

AOB:

- Reminder of GGUS ticket on BDII issue - Jon: checked this morning and out of date for close to 80h 54803.
- Notes from last USAG meeting sent out earlier. Please read <http://indico.cern.ch/getFile.py/access?resId=0&materialId=minutes&confId=81363> . Next meeting - Doodle setup <http://doodle.com/uuuyhysavtptyc52g> please choose. Will discuss issues concerning all VOs and sites.

- Issue of recalculation of availability. Follow-up?

Wednesday

Attendance: local(Harry, Eva, Jan, Jamie, Jean-Philippe, Nicolo, Timur, Roberto, Patricia, Antonio, Steve, MariaDZ);remote(Jon, Onno, Michael, Joel, Angela, Rob, Rolf, John Kelly, Jason, Alessandro Italiano (INFN-T1))).

Experiments round table:

- ATLAS reports -
 - ◆ FTS channel IN2P3-SARA stuck: FTS channel monitoring did not show any Active FTS jobs, there were only Ready and Finished jobs. Problem under investigation, "FTS racing" with LHCb excluded so far. GGUS:55008 [↗](#). Similar issue observed during weekend on BNL-CNAF FTS channel (cause not found), GGUS:54943 [↗](#).
 - ◆ RAL on scheduled downtime.
- CMS reports - Highlights:
 - ◆ CERN
 1. Errors 'DESTINATION error during TRANSFER_PREPARATION phase: [SECURITY_ERROR] Current user does not own this request!' in imports from all sites to CERN yesterday 2010-01-26 17:30 UTC, recovered in 1-2 hours.
 - ◆ CERN CAF
 1. 4 files from the most recent collision skim have wrong checksum on cmscaf - GGUS TEAM ticket submitted.
 - ◆ T1s
 1. Issue with job status monitoring on CREAM CE identified by developers.
 - ◆ CCIN2P3
 1. Some instability in SAM and JobRobot after end of scheduled downtime (probably for additional batch node maintenance), now OK. [Rolf - when did you observe these problems? Nico - will check elog and add timestamp to minutes]
 - ◆ FNAL
 1. Timeouts in FNAL-->T2 exports
 - ◆ CNAF-FNAL
 1. Timeouts in transfers FNAL-->CNAF
 - ◆ ASGC
 1. Following up on deployment of SLC5 software releases
 - ◆ T2 sites issues:
 - ◇ T2_ES_IFCA - problem in job reporting to Dashboard from production jobs
 - ◇ T2_FI_HIP - progress in software installation.
 - ◇ T2_PK_NCP - CERN-NCP channel demonstrated good stability, requested increase in number of files on FTS channel on FTS-T2-SERVICE to saturate site bandwidth.
- ALICE reports [↗](#) -
 - ◆ T0 site o CASTOR update announced last week has been completed this morning. No issues seen by ALICE following upgrade. o ALICE Central services: Installation of the latest ALIEN version in the central services has just been done. The main motivation is to get the latest version of xrdep o Issue with access to voalice14 solved just before meeting.
 - ◆ T1 sites o No issues to report
 - ◆ T2 sites o Hiroshima: Local setup of the vobox finished. Ongoing with the LDAP configuration to put it in production (hopefully today...) Request for new VOBOX for Cape Town - once finished will enter production.
- LHCb reports - System is empty apart from users jobs.

- ◆ T0 sites issues: o Asking about the status of the second machine behind LFC-RO. JPB - no progress. [Steve will followup] o Question on VOBOX - pending since one week.
- ◆ T1 sites issues o RAL : Downtime o CNAF:CREAM CE still observed problems if using sgm account. (GGUS opened on Monday)
- ◆ T2 sites issues: o none

Sites / Services round table:

- FNAL - 1) one of BDIIs at CERN was taken out and then later another failed in same mode. Reported. This morning update that new code applied and now seems ok. Thanks! 2) Short robotic failure yesterday - worked as expected, everything queued. Source of some complaints about slow transfers - some source files on tape and stalled for a few hours. 3) Rolling upgrade of dCache almost finished 4) FNAL - CNAF some due to robot outage, rest due to FTS server at CNAF & haven't looked any further. Alessandro (CNAF) don't have any info on this - investigating the ticket opened by LHCb on CREAM CE... Nico - checked yesterday, at FTS level things seem ok but underlying transfers are flow. Maybe SRM?
- NL-T1 - problems with FTS channel IN2P3-SARA under investigation. Tomorrow scheduled maintenance on SARA SRM to activate new kernel.
- BNL - ntr
- OSG - ntr
- IN2P3 - ntr
- KIT - ntr
- RAL - other than ongoing downtime ntr
- ASGC - CMS ops problems confirm job will take more than 7 hours. :Probably reinstall another fileserver.
- CNAF -

- CERN - CASTOR - had SRM ATLAS hiccup again today. Prob concurrent activity on scratch pool. CMS reported corrupted checksums. Systematic? If so turn off.. Noticed on LHCb data pool some fairly large (75GB files). Mess up allocation logic - will send list. Joel - this is normal activity. Jan- need to discuss. Risk running out of disk space. CASTOR ALICE Upgrade to 2.1.9-3 completed as mentioned in ALICE report.

Release report: deployment status wiki page - last update for this meeting(!) Update to gLite 3.2 in staged rollout. Relevant content is new ARGUS package - new authorization service for SL5. Porting of DPM to SL5. glxec tests to gLite SWAT client on WNs. AFAWK only way now to understand which sites run glxec and which version. Have to be deployed on WNs to be effective. Release will be out be end this week or beginning next. Late in corresponding update of gLite 3.1. Dependency of dCache/dcap clients. This update will contain FTS 2.2 version...

AOB:

Thursday

Attendance: local(Timur, Julia, Jan, Ignacio, Jacek, Steve, Maria, Jamie, Miguel, Harry, Roberto, Jean-Philippe, Stephane, Lola, Nicolo);remote(Michael, Ronald, Gang, Angela, Jon, Rolf, Jason, Gareth, Jeremy, Brian).

Experiments round table:

- ATLAS reports -
 1. RAL downtime for Castor
 2. Announcement: Change of FTS setup for the number of parallel transfers to better match the T1 share. This will be tested on monday 1st Feb. A bunch of transfer will occur each 5

minutes :

- ◇ RAW+ESD sent to T1 sites according to T1 shares (N): 175 GB * N
- ◇ AOD+DPD sent to all T1 : 65 GB to each T1

Brian - RAL would like to set a maximum number. Stephane - ok will communicate, maybe need 2 tests, Simone will send mail.

- CMS reports - Highlights:
 - ◆ T0
 - 1. Drop of 70 TB on T0EXPORT pool, queued transfers Remedy #CT0000000657302
 - ◆ CERN CAF
 - 1. Wrong checksum on cmscaf confirmed to be due to 2.1.9-3 gridftp checksummin bug.
 - ◆ T1s
 - 1. Issue with job status monitoring on CREAM CE identified by developers.
 - ◆ CNAF
 - 1. JobRobot failures at CNAF caused by CREAM CE reconfiguration bug. Excluding CE from resources monitored by CMS.
 - ◆ CNAF-FNAL
 - 1. Quality improved but still some timeouts in transfers FNAL-->CNAF
 - ◆ ASGC
 - 1. Following up on deployment of SLC5 software releases
 - ◆ T2 sites issues:
 - ◇ T2_ES_IFCA - problem in job reporting to Dashboard from production jobs solved - jobs were being killed without reporting by hard limit in stack memory.
 - ◇ T2_FI_HIP - SLC5 deployed.
 - ◇ T2_PK_NCP - CERN-NCP channel demonstrated good stability, requested increase in number of files on FTS channel on FTS-T2-SERVICE to saturate site bandwidth.
 - ◇ T2_PT_NCG_Lisbon: SAM test failures for StoRM backend crash.

Ignacio - T)EXPORT issue is due to error when preparing draining machines to be replaced by new boxes (machines going out of warranty). Machines were not in disk pool.

- ALICE reports [↗](#) - GENERAL INFORMATION: Changes in the MC cycle today. instabilities in the job profile are expected. Production is almost stopped at this moment.
 - ◆ T0 + T2s: ntr
 - ◆ T1 sites - CNAF: The submission to the LCG-CE backend is stopped. Wrong information provided by voview for the ALICE queues at the site is causing an overload of agents submitted to the site. Site admin aware of the problem and following the issue with him
- LHCb reports - Relaunched the L0HLT1 stripping for b and c inclusive MC09 events with a "corrected" workflow description (less input files for avoiding the enormous output files reported yesterday) [Will eventually delete the v large files]
 - ◆ T0 sites issues:
 - ◇ Looking for news from Steve about RT CT643684 (vobox replacing old Log SE service) and CT654872 (extra FE behind LFC-RO) (Jamie - I also mentioned this to Maite for her to follow-up.)
 - ◇ Verified that 75 GB files reported yesterday by Jan weren't to be considered "normality" but rather a mistake on a stripping production definition and then to be scrapped.
 - ◆ T1 sites issues: RAL : Downtime.

JPB - LFC being installed now. Steve - VOBOX: confusion - in progress.

Sites / Services round table:

- BNL - firmware upgrade of core switch completed yesterday; problem observed earlier should be fixed. Q for Stephane - all data to disk or also tape? A: disk but will check. Will start Monday morning CERN time.
- NL-T1 - NIKHEF I/O errors on a disk server. Expect result in data loss for ATLAS! Investigation. 1 d/s unreachable, 3 others in same rack r/o mode to protect data, escalate to vendor. Stephane: which end-point? A: NIKHEF. TBN18.
- ASGC -
- KIT - ntr
- FNAL - scheduled network outage to reboot switch. Affected CERN-FNAL traffic, failed over as planned. Q: mailing list? Please add me!
- IN2P3 - ntr
- OSG - ntr
- RAL - in middle of fairly significant outage. Big work. One small thing - 1 of ALICE VO boxes failed to reboot - rebuilt. Restoring ORacle DBs onto disk arrays. LFC & FTS DBs work went fine and services back in prod. Some problems with CASTOR DBs hence extended end of outage from end of today until tomorrow pm. Some detailed config issues, since problem on one of Oracle RACs. Since migration Oracle can't see data. Call open with Oracle and contact with CERN team - trying to resolve.
- GRIDPP - ntr

- CERN DB - progressing with latest security patches. Both online & offline LHCb today; all other production DBs next week.
- CERN DB CASTOR - downtime for all CASTOR instances at CERN. Next Wednesday 09:00 - 10:30? Please confirm if ok... Jacek - why not rolling? A: don't know. Jamie - please provide an explanation for this meeting and the experiments!

AOB:

Friday

Attendance: local(Jamie, Steve, Harry, Miguel, Gang, Maria, Nilo, Jacek, Simone, Jean-Philippe, Timur, Nicolò, Patricia, Jan, Roberto);remote(Xavier, Jon, Gonzalo, Micheal, Onno, Stephane, Rolf, anon, Jeremy, John, Tore Mause (NDGF), Jason, Rob, Luca).

Experiments round table:

- ATLAS reports -
 1. FTS transfers CNAF-> NAPOLI stuck during 12 hours (GGUS:55095 + Savannah 62049)
 2. Typo in DDM config file introduced last night. Affected all ATLAS Grid activities during the night until correction this morning. For example, explains the hole with MC production.
 3. RAL+FZK missing for monday throuput test (loose 20% of T1) - have to reassign b/w to other sites and check overall bandwidth.

Simone - no chance to move throughput test. Activities for rest of week - date fixed! Can have dedicated test of FZK + RAL the week after. Michael - good opportunity to test Use Case where part of T1 capacity is not available. Send to "sister site" - BNL volunteer to accept more! Simone - for data assigned to site reshared to other sites according to share. BNL will get 25% + 25% of what RAL+FZK cannot take. Michael - could be more! Miguel - skewing distribution? Simone - if reshare this is ok. CNAF requested data to tape too to test StoRM + TSM. Will be done for this site. Can publish export rates. (Y!)

- CMS reports - highlights:
 - ◆ FNAL-IN2P3:

- 1. Transfers not starting - sites contacts notified
- ◆ FNAL-CNAF
 - 1. Ongoing investigations on FNAL-->CNAF timeouts [timeout has been increased for the time being whilst being understood]
- ◆ ASGC
 - 1. Following up on deployment of SLC5 software releases
- ◆ T2 sites issues:
 - ◇ T2_RU_IHEP, T2_EE_Estonia - SAM test errors
 - ◇ T2_BR_UERJ - unscheduled downtime
 - ◇ T2_PK_NCP - CERN-NCP channel demonstrated good stability, requested increase in number of files on FTS channel on FTS-T2-SERVICE to saturate site bandwidth.
- ALICE reports [↗](#) - GENERAL INFORMATION: Following the conclusions of yesterday's ALICE TF meeting, beginning the deprecation of the LCG-CE (and gLite-WMS) services at those sites providing already a stable CREAM-CE for ALICE. In all those cases where the site is providing a 2nd VOBOX, this node will be implemented in LDAP as a backup service of the production VOBOX
 - ◆ T0 site - Shutdown of CASTOR next week for the implementation of a security patch: 1h30 of downtime easily to absorb. No major objections for the intervention (as usual, a "go back" solution in case something goes really wrong is needed), this for both Tuesday and Wednesday. During the next couple of weeks, if anything has to be done ALICE prefers it to be done in the morning, when activities are - in general - less data-centric.
- LHCb reports - Stripping on MC09 b and c inclusive events + normal users activity: less than 1000 jobs
 - ◆ T0 sites issues:
 - ◇ Noticed yesterday at around 14:00 all SAM tests against our LFC at T1's failing systematically. This is certainly in line with the scheduled intervention for the Oracle security patch on the LHCBR offline (13:30 -15:00). Announced to be rolling fashion the DBA envisaged indeed some perturbations in the connection during the 90 minutes of intervention [Jacek all applications should be prepared to reconnect - LFC does!]
 - ◇ LHCb has no problem to run the CASTOR intervention next week for the Oracle patch. Brought the question whether these jobs already running in the batch system will be frozen or just let dieing/hanging. (input for a more general point about procedures for the Thursday meeting)
 - ◇ vobox and extra LFC node. Looking for status.
 - ◆ T1 sites issues: IN2p3: the announced stress test against new client of gsidcap has to be postponed (the suite from DIRAC is not ready yet). Unlikely to be ready for the next week due to the software week .
 - ◆ Problems with services: We see a lot of tests/probes non-lhcb specific polluting our production services: This is the content of a directory under the global space of lhcb LFC. We kindly ask responsible (Nagios managers) to agree with LHCb about the convention used to put everything under a test area also available at the SEs. [Details in full report] Steve - quite a few ticket updates for VO box. JBP - other LFC node is under test - being put in to load-balanced DNS. Monday? Harry - delivery of 65 palettes (1st of 3 deliveries)

Sites / Services round table:

- KIT - ntr
- FNAL - yesterday continued to investigate FNAL-CNAF transfers. Default is using CNAF FTS server. Using FNAL FTS server works ok. Believe timeouts at CNAF have to be tuned. Looking at FTS - PhEDEx issues for IN2P3 and also Florida.
- PIC - ntr
- BNL - ntr

- NL-T1 - announcement: SARA CE and CREAM CE are in scheduled maintenance on Tuesday - to be moved to different h/w.
- IN2P3 - ntr
- ASGC - ntr
- GRIDPP - ntr
- RAL - having problems with DB migration. 2 identical EMC disk packs. One works and one doesn't. Intermediate plan to bring up one of disk packs running ATLAS and LHCb - up by Monday lunch. Others by Tuesday lunch. All "fingers in the air" but this is the plan...
- NDGF -
- CNAF - problems with transfers from FNAL. news from Jon are interesting. Jon - you are running slightly older FTS server than us. Older server doesn't expose all timeout values so we can't check how its configured. Luca - need to put FTS manager in contact with yours. We weren't sure if timeout increase was the right approach. Why does this problem occur only since a couple of days? Testing StoRM for ATLAS with tape b/e. Hopefully will switch also for ATLAS next week with 1 day stop of service. Simone - after throughput and reprocessing tests!
- OSG - we had some problems with ticket exchange GGUS - OSG. Problem was from OSG->GGUS. Bug fix applied Wednesday. Issues with updates into GGUS fixed. Another broker outage and resending records - 4th time this month! OSG will be asking for recalculation again ;-(Jon - my understanding is that updates and fixes would be automatically taken no matter how big or small. FOLLOW-UP

- CERN - confirm date of CASTOR update for Tuesday morning. Roberto - what will happen to batch queues? Miguel - we will stop scheduling new batch jobs and will pause special queues on request. Many jobs don't handle pausing well - some do no / little I/O. Maria - is it clear why security patch cannot be rolling? Nilo - it is other patches than need to be applied and require full access to database. Jamie - use this to "kick-off" Risk Analysis procedure as discussed at January GDB. Motivation | Risk Analysis | Post-Mortem

AOB:

-- JamieShiers - 22-Jan-2010

This topic: LCG > WLCGDailyMeetingsWeek100125

Topic revision: r10 - 2010-01-29 - JonBakken



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