

Table of Contents

Week of 100201.....	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.....	5
Thursday.....	8
Friday.....	10

Week of 100201

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(Patricia, Jamie, Steve, Miguel, David, Gavin, Nicolo, Alessandro, MariaDZ, Jean-Philippe, Timur, Simone, Stephane, Harry, Lola, Roberto, Eva, Ueda);remote(Jon, Xavier, Rolf, Gang, Brian, Gareth, Angela, anon, Daniele).

Experiments round table:

- ATLAS reports - Summary of problems then throughput test later...
 1. Central catalog : logging issue preventing DQ2 working (affecting production/analysis). At ~3:00 am on saturday, /tmp got full for the 3 Central Catalog machines. ADC Expert called in

the morning by ATLAS. Misunderstanding between sudo access and interactive access to clean the /tmp/env.txt (has been clarified in the meantime). Quick answer from Flavio when mail sent to Central Service team. Problem solved around 12:00.

2. Panda task dispatcher (Bamboo) had problems to interact with Atlas Central Database (starting on friday afternoon but quickly hidden by Central Catalog problem). DBA support was contacted on saturday afternoon and problem was solved on saturday late afternoon. It affected the MC production
3. Sites issues o JINR-LCG2 : Problem of certificate during few hours (GGUS : 55120) o LIP-LISBON : No more access to SE (GGUS : 55118)

Simone started this morning. Smooth -but 5% of CERN-CERN transfers have CASTOR error (too many threads). BNL went down about 11:00 but solved since 20' ago. Michael - dCache problem, namespace non-responsive? Not fully understood... Transfer rate ~10Gbit/s - one of 2 links completely saturated... (Stephane) - throughput tests continues until this evening.

Miguel - ATLAS T0 people reported problem with migrations. Still under investigation. Daemon which is supposed to queue migrations doesn't appear to be doing it fast enough. Brian - too many threads is something we at RAL would also like to follow as we've seen this here (RAL). Any fixes - please communicate!

- CMS reports -
 - ◆ Dashboard
 1. SAM test results unavailable on old page, accessible with new page
 - ◆ T1 rereconstruction and skimming in progress
 1. PIC issues with file staging - possibly problem with tape staging protection GGUS #55121
 - ◆ ASGC
 1. Following up on deployment of SLC5 software releases
 - ◆ T2s:
 - ◇ Some MC jobs crashing at T2_US_Florida for memory limit
 - ◇ T2_ES_IFCA, T2_UK_SGrid_RALPP, T2_US_MIT, T2_PT_LIP_Lisbon, T2_FR_IPHC SAM test errors
- ALICE reports [↗](#) - GENERAL INFO: The startup of a new MC cycle is foreseen for this weekend after a weekend with a small production. The central activity of the weekend has been a pass 3 reconstruction of raw data at FZK. In addition stress tests of the gridftp server installed at the gLite3.2 VOBOXES have been performed with the setup of Subatech (ALICE T2 in France)
 - ◆ T0: No issues to report in terms of services. There are still 2 pending VOBOX registrations sent to PX support by the 22nd and the 27th of January:
 - ◇ CT656042: Registration of a VOBOX in Clermont
 - ◇ CT657018: Registration of a VOBOX in CapeTow(n)
 - ◆ T2s: Stress tests at Subatech of the gridftp server installed at the gLite3.2 VOBOX have succeeded. Thanks to the support provided by the site admin and the CREAM-CE developers
- LHCb reports - Software week
 - ◆ T0: CASTOR intervention tomorrow: LHCb asked to suspend running jobs on its dedicated queue (grid and not-grid)
 - ◆ T1s: IN2P3: 1h. unscheduled downtime on Saturday but received the notification from CIC portal 13 hours after the END (see pics [in full report])
 - ◆ T2s: ITEP: wrong mapping of FQAN /lhcb/Role=production. The issue there is more general and affects the backup mapping solution based on old static gridmap-files. LHCb clearly states in its VO Id Card that static mapping through gridmap file should only maps users to normal pool account (.lhcb) instead of to any other super-privileged account.

Rolf - message from CIC portal arrived at about 12:30 - this was due to fact NREN established a peering point

in Paris which was needed. Unscheduled announcements can be announced "after the fact" to keep you informed.

Sites / Services round table:

- FNAL - tomorrow morning security update to Oracle FTS instance.
- KIT - today started downtime for ATLAS dCache for migration pnf's to Chimera. Smooth so far... Updating WNs since Thursday last week. New kernel + reboot. Some went back into production too early and hence some SAM tests might have failed - problem understood and fixed.
- OSG - some tickets that we are waiting on updates for - a couple of ATLAS tickets. 54717, 54454. See under AOB that Maria has added some stuff. Discussion on 1st ticket - not much to report. On 2nd ticket closed in November and now re-opened. Maria - was this because it wasn't solved? Kyle - closed as there were multiple tickets for the same issue. GIP 1.1.8 update came out recently and should fix problem - will double check Burt.
- BNL - nothing extra
- IN2P3 - issue with SRM during w/e but hasn't been noticed by LHC experiments. Had to restart service Saturday evening. Had power cut a few minutes ago - problem stopped many WNs. Several jobs which crashed - more to come.
- ASGC - SIR of power surges has been uploaded. First stage of protection for oracle and CASTOR services done last Friday. 2nd stage of additional protection for WLCG services will be completed this month.
- RAL - still have CASTOR (DB issue!) and batch down (batch because of CASTOR). CASTOR outage until Wednesday. Migrated DB back on to original disk arrays before last October. Power supply from UPS noisy and led to instabilities. Migrated DBs back - resilience from SAN underneath has not be there. Many problems in getting resilience back - still not fully there. Testing CASTOR on the system with aim of bringing it back. Hopeful before current outage over (Wed lunch). Independently planned migration of 3D DBs onto another instance of the storage. Going on but slowly.. Site network work on Tuesday 9th Feb.
- CERN - discussed change record and risk analysis for CASTOR + SRM + DB: nameserver tomorrow, srm ATLAS Wednesday.
- CERN DB - interruption on ATLAS AMI replication during w/e due to network problems. Continue deployment of latest security patches. Problem reported by LFC replication was not a real problem of 1h downtime - some replication delays due to intervention. Tests are repeated every hour which gives this granularity in "down time" report.

AOB: (MariaDZ) On OSG https://gus.fzk.de/ws/ticket_info.php?ticket=55115 has a GGUS-OSG synchronisation problem already being discussed by the developers. Today's escalation reports show https://gus.fzk.de/ws/ticket_info.php?ticket=52982 opened as urgent last November. Please put an update in the ticket, not in this meeting's notes. Burt Holzman (CMS) was working on this.

Due to IN2P3 / CIC portal issues broadcasts which caused knock-on problems. Being investigated.

Tuesday:

Attendance: local(Dirk, Eva, Miguel, Jamie, Maria, Nicolo, Timur, Jean-Philippe, Julia, Nilo, Ueda, Ewan, Stephane, Simone, Alessandro, Roberto, Patricia, Andrea, MariaDZ);remote(Jon, Michael, Rolf, Angela, Jens, Jeremy, Gareth, Rob).

Experiments round table:

- ATLAS reports -

Tuesday:

1. Restart transfers to RAL (under validation) [Discussion this morning, start func. tests end morning, since ok restarted almost all transfers - looks promising - need full day of tests.]
 2. Results of the Throughput Test: full post mortem during the ATLAS Jamboree (Tuesday Feb 9th). It is under investigation backlog in PIC and TRIUMF after end of the test (plot attached to full report). Sites have been contacted. [Yesterday 9am to 9pm, no more new data by midnight. 2 sites - TRIUMF and PIC - had backlog at end. Took around 6h to complete. Rate to these sites did not increase after others had completed so not congestion at source. 15 actives transfers - no change - to PIC. Not understood. Will repeat for these two sites - started with PIC at 14:00 shipping fake raw files of 5GB each. ~5MB/s per file over OPN. x 15 active transfers gives rate seen yesterday. Sent info to PIC. To be verified if true for TRIUMF, possibly starting today.
- CMS reports -
 - ◆ PIC
 1. issues with file staging for exports fixed, files needed for reprocessing manually prestaged
 - ◆ ASGC
 1. New software area configured, CMSSW for SLC5 deployed
 - ◆ CNAF
 1. lcgadmin SAM test jobs not running/timing out on some CEs due to long running jobs in the queue (behind some ATLAS jobs!) Still 1 CE ok so overall availability OK.
 - ◆ IN2P3
 1. Batch system issues
 - ◆ T2 issues: Ongoing SRM SAM test failures at T2_US_MIT (unsched down since power failure)
 - CMS weekly update
 - ◆ T0 - global run Wed/Thu, next week data taking from Monday (continuous); possibly also replay of existing data with latest CMSSW version.
 - ◆ T1 - no new requests; continue tests, backfill, issues reported above
 - ◆ T2 - ongoing MC; rerun of some samples effected by generator file event duplication - output needs to be regenerated.
 - ◆ Reminder of WLCG procedures regarding sched/unsched downtimes. (Extension is considered unsched for ex.)
 - ◆ FTS configuration - tune channel parameters if needed: Start with CERN then T1s. First report in CMS next Monday.
 - ◆ SL5 WM migration - 4-5 sites still to do
 - ◆ ASGC - discussions on tape family definitions for 2010.
 - ALICE reports [↗](#) - GENERAL INFORMATION: As we announced yesterday, the new MC cycle has been started yesterday afternoon (over 10K concurrent jobs)
 - ◆ T0 site o Waiting for ALICE feedback after the CASTOR operations scheduled for this morning
 - ◆ T1 sites o CNAF: GGUS ticket 55156 submitted this morning to track the problem observed in the information provided by voview for the ALICE dedicated queues. The wrong information provided by the info system might provoke an overload of submitted jobs. ALICE has stopped the submission to the LCG-CE of the site until the issue is solved.
 - ◆ T2 sites o Huge amount of vbox-proxy processes running inside the vbox in Poznan. The problem was announced yesterday, all processes were stopped and old vbox-proxy processes manually killed. The startup of the services was done yesterday afternoon and the same issue has been announced this morning. Studying the issue together with the site admins
 - LHCb reports - No activity. Software week. Under discussion within the collaboration: to extend to

T2's (under specific and restrictive conditions) the possibility to host distributed analysis (besides T1 and CERN-CAF). Interesting this talk with proposed amendments to the LHCb Computing Model (to be approved by CB).

- ◆ T0 sites issues:
 - ◇ Got another machine behind LFC_RO at CERN. Everything seems to be OK 😊
 - ◇ CASTOR intervention took longer than expected (half an hour more than expected).
- ◆ T1 sites issues:
 - ◇ RAL: LFC SAM tests for streams replication are failing systematically since yesterday. Most likely due to the 3D intervention over there.
 - ◇ IN2P3: perturbation with the new MySQL DB put behind BQS last Tuesday causing some jobs not being submitted through.

Sites / Services round table:

- FNAL - Oracle security ongoing. Should finish within 2 hours.
- BNL - ntr
- IN2P3 - during the call yesterday suffered from power cut. After some minutes we had to stop WNs. Had to stop all WNs in about 20'. Jobs which were running at that moment crashed. Power came back later - reason lies with supplier - broken cable somewhere in the power network. Switched to other branch - took some time. Powered up WNs during pm and night. Except WNs for SL4 - came back this morning. Batch system had upgrade to new h/w and another version of MySQL etc. Meant to speed up everything but had opposite effect 😞 Not sure of reasons - h/w is faster but new versions slowed everything down. Probably rollback in some hours.
- KIT - ntr
- ASGC - concerning CMS tape repacking phase 1 finished. 200TB freed. Phase 2 ongoing. Reconfigure tape families during this time too.
- NDGF - short intervention of dCache servers tomorrow for minor upgrade. 'fairly transparent'.
- GridPP - ntr
- RAL - services have started coming back. CASTOR restored about 11:00 local time - running ok since. FTS restarted with channels low but now increased. Started batch up now. 3D service - at disk and short outage yesterday - largely done OK. Streams turned off but re-enabled during this meeting.
- NL-T1 - downtime on CE and batch. Finished and back in production. NIKHEF - disk server serviced by vendor. No problems nor indications of data loss. Hope fully back in prod in a few days.
- OSG - 36h of SAM records asked to resend - done. recalculation necessary?
- Dashboards - problem reported by CMS - SAM portal problem - fixed.
- DB - intervention in ALICE pit now over. Online DB switched over to primary at the pit; standby in CC.
- Nameserver upgrade this morning. 30' delay but nothing went wrong - took more time than when tested. At 2.1.8.3. Nilo - no need to use special scripts so completed ahead of time.

AOB: (MariaDZ) GGUS Release tomorrow 2010-02-03. Tier1s please remember a test alarm will be initiated by the GGUS developers in the afternoon and after every release as per <https://savannah.cern.ch/support/?111475#comment13> 📧 The alarm notification email will be signed by the GGUS certificate as always.

Wednesday

Attendance: local(TIMur, Przemyslaw, Lola, Luca, Ale, Jamie, Maria, Simone, Roberto, Jean-Philippe, Ueda, Miguel, Harry, MariaDZ, Patricia);remote(Jon, Joel, Rolf, Pepe, Gonzalo, Michae, Onno, John, Jeremy, Gang, Jens, Angela, Rob).

Experiments round table:

- ATLAS reports -
 - ◆ There was a glitch in accessing ATLAS Computing Operations ELOG before noon. [Ale - only 1 person handling elog: need to find a solution]
 - ◆ Results of the Throughput Test: see plots in full report.
 - ◆ Tests have been repeated for PIC, starting from 2PM CET. The problem has been identified in the transfer rate per single file, hardly exceeding 10MB/s (times 15 active transfers gives the observed 150 MB/s of throughput). Explanation from PIC: "As we changed to gridftp2 the transfers are going to the pools directly and not through the gridftp doors, seems windows size is not well negotiated or tuned at PIC pools (128kb now), in the past gridftp doors were well tuned for this, so this effect is new. Xavi Espinal". So during the test the parameters have been tuned to window size 256k, max windows size 4MB. The transfer rates increased to 750MB/s immediately. Problem solved. * Tests repeated for TRIUMF starting from 9PM CET. The problem comes from the only three pools with SUN hardware (same as PIC) which deliver ~6MB/sec/transfer, while ~50MB/sec/transfer on others. Some tuning has been tried for the SUB boxes but with no appreciable result. More tuning is needed.

Simone: PIC started 2pm yesterday and after <1 h released source of problem. Debugging for TRIUMF started much later. Tried same trick as PIC but didn't help - in contact with dCache expert in TRIUMF who should do more tuning. If you hit one of the above 3 pools you get slow transfers. Maria - 50MB/s is a target?
 A: yes; problem identified but not solved.

- CMS reports -
 - ◆ PIC
 1. Local tests ongoing to check why the service certificate used by central DataOps team fails to stage files. Apparently, everything is well setup. dCache developers having a look as well.
 - ◆ ASGC
 1. New software area configured, CMSSW for SLC5 deployed and all all slc4 production releases have been reinstalled on the new CE. Action closed.
 - ◆ IN2P3
 1. Batch system issues; reverting to previous BQS version this morning. Leaving the upgrade to newer version to end February.
 2. CCIN2P3->SgridRALPP transfer errors. Seems SgridRALPP endpoint is not on CCIN2P3 FTS service.xml file.
 - ◆ T2: Ongoing SRM SAM test failures at T2_US_MIT

Jon - missing file at FNAL. Don't see any open ticket and not aware of any missing files. Pepe: Savannah #112498 will reassign.

- ALICE reports [↗](#) -
 - ◆ T0 site - ntr
 - ◆ T1 sites
 - ◇ CNAF: Issue reported yesterday and included in the GGUS ticket: 55156 has been solved and verified by Alice this morning. The issue was concerning the wrong information provided by voview regarding the Alice queues of the sites
 - ◇ RAL: the site came back after an outage today. Manual operations for the startup of the AliEn services have been needed. The site is back in production
 - ◆ T2 sites
 - ◇ Prague T2: Today in the afternoon, the Prague batch system will be off due to the intervention on the torque server. The intervention will take about 2 hours. Jobs running on the farm at that time will crash.
- LHCb reports - no production on-going.
 - ◆ T0 sites issues: o Request to verify whether the master instance of LFC at CERN has new

delivered VOBOXes at CERN in the list of trusted hosts. o problem of synchronisation between VOMS and VOMRS ()

◆ T1 sites issues:

◇ IN2p3: Banned because of the intervention on BQS backend. Received announcement 1h after intervention has started. Why the delay in receiving the announcement? Rolf - don't know message arrived late. One reason might be that downtime was declared late - will check. MariaDZ: please don't send mail to Steve or anyone else but open a GGUS ticket! Joel - then close voms*-support mailing lists?
A: Y

◆ T2 sites issues:

◇ Shared area issue at Manchester.

Sites / Services round table:

- FNAL - looked at ticket and file is available - transferred successfully! Noticed yesterday that FTS does not properly cleanup - gathering info and will submit ticket. JPB- will check and fix.
- IN2P3: One direct update on 1h delay - times in announcement are in UTC! Downtime was announced at 08:00 UTC = 09:00 local. Joel - last one was sent at 12:18 UTC arrived at 14:10 localtime. Rolf - batch system issue: did some serious testing on new versions. About 50k jobs loading new machines - before putting new version & m/cs in production. First signs of performance degradation came up ~4 days after production. Not able to detect real reason. Depends on time between installation and some days later - suspect DB problem? Not able to confirm. Simply needs several days to show up - not enough time to reproduce. Hence revert. In parallel will leave new version on a test system and try to understand and fix. No degradation seen when testing for several days but not under real load. Seen with 50K jobs / day over ~ 4 days.
- PIC - ntr
- BNL - slowness in CondorG submissions - investigating.
- NL-T1 - ntr
- RAL - brought CASTOR and batch farm back after outage. Gareth is preparing a post-mortem.
- GridPP -
- ASGC - yesterday CMS SAM CE installation tests failed 00:00 - 21:00. CMS SW team installing packages during this period. Ticket closed. Next Wed/Thu will have power-off for 23h for regular safety check at Academica Sineca. ASGC also affected - will confirm date.
- NDGF - ntr
- KIT - downtime for ATLAS due to Chimera migration is finished. All tests successful - plan to go back to production tomorrow morning. i.e. ahead of schedule. Simone - should we wait for green light to start func. tests? Maybe from tomorrow morning? Angela - func tests can start now. Will end downtime and then you can start production.

* CERN DB - replication to BNL for ATLAS. A few tables (~3) have some missing rows - being investigated. Goes back a few weeks - will require instantiation of a few tables. Trying to figure out why...

AOB:

- FTS patch for delegation has been certified. CERN is running a non-certified version of FTS. See no reason to run 2.2.2 - not an official release - go to 2.2.3 asap. Will prepare information for sites by next week's Tier1 Service Coordination meeting.
- GGUS release today. One of the new feature is the possibility to open a ticket on behalf of a 3rd party. To facilitate incident reporting via GGUS a new button is implemented on the ticket submission form for label "Notification Mode" with value "Never". This will allow anyone with a valid certificate from a trusted CA to open a GGUS ticket on behalf of a user, without having to receive notifications. The user should be put in Cc and will be informed of all progress and the solution.
<https://savannah.cern.ch/support/index.php?111183>

- GGUS test alarm tickets to CERN. Contained no info so operator can't route them. Test that ticket gets to CERN? If more then we need more info... MariaDZ: this is correct - test of delivery to site.

Thursday

Attendance: local(Ueda, Miguel, Harry(chair), Roberto, Gang, Timur, Lola, Gavin, Ewan, Cerdic, Steve, Ignacio);remote(Jon(FNAL), Brian+Gareth(RAL), Angela(KIT), Michael(BNL), Pepe(CMS+PIC), Ronald(NL-T1), Jeremy(GridPP), Rolf(IN2P3), Gunter(GGUS)).

Experiments round table:

- ATLAS reports - 1) Several monitoring on SLS were affected yesterday evening and this morning by the deployment Oracle security patches on lemonops DB. It was announced on ssb, but was difficult to notice and we did not notice. Also the SLS of lemonops DB (http://sls.cern.ch/sls/service.php?id=DB_lemonops) showed 100% availability all the time. The incident this morning seems to be a separate issue according to the remedy ticket (CT659140). 2) There was a problem in transfers to RAL for a while this morning but the issue has been solved (ggus:55287). 3) FZK has been put back into production. 4) Some corrupted files have been found at NDGF (first spotted in Lyon) which is a reason for wanting to move to FTS 2.2 with its checksumming functionality.
- CMS reports - 1) PIC: Service certificate used by central DataOps team failed to stage files as there was a mis-configuration at dcache level ([[#5419](http://www.dcache.org)] Tape protection configuration reload). In the current release, one has to specify an special config entry ["*" "/cms(/.*)"] to catch all CMS roles. Patch applied at PIC and it's under test at this moment. 2) Would like a site report from ASGC next Monday prior to them being added back into CMS production. 3) Reduced capacity in the CERN T0 export CASTOR pool has now been restored. 4) Timeout errors on FNAL to CNAF transfers have been fixed but a ticket is open to follow up with transfer rate tests.
- ALICE reports - 1) GENERAL INFORMATION: MC production ongoing with small number of jobs plus analysis train activities. 2) T1 sites: RAL Issues to start the alien services with the local expert credentials. Following the issue with the site admin. 3) T2 sites: no major issues to report in terms of the T2 behavior. CREAM-CE setup at CapeTown: problem with the authentication. Following the issue with the site admin. This is the last step before putting the site in production
- LHCb reports - 1) T0 sites issues: New delivered VOBOXes have been registered in the list of trusted hosts by LFC master. To be effective this change has to be propagated to the Quattor template - expected to be done next week. 2) Lancaster T2 issue with the mount point of SL5 sub-cluster shared area. They are using one single endpoint for the CE service pointing to two different sub-clusters with different OS. 3) UKI-LT2-IC wrong mapping of Role=production (again the issue of a too generous mapping in case of gridmap-file mechanism is used). 4) New GGUS portal: TEAM tickets lose the information about concerned VO resulting to affect "none" instead of the expected "lhcb". Later reply from G.Grein: There was a typo in the (GGUS) code. This is already fixed. 5) VOMS <-> VOMRS synchronization: probably due to the original UK CA certificate a user has been firstly registered (and expired long while ago). Steve fixed it by hand with Tania's help and reported it was an Oracle bug that happens sporadically at any site. Many other users potentially affected by this.

Sites / Services round table:

- FNAL: 1) Noticed that one of the bdi servers at CERN was out of sync at FNAL by 10 hours. GGUS 55262 was opened but it was back to normal when checked later. 2) A ticket has been opened for the problem reported yesterday of FTS not cleaning up /var/log/pid files which makes startup problematic. 3) Network experts at Caltech/CERN/FNAL/BNL have been seeing huge network speed oscillations. Average rates are good e.g. at 600 MB/sec from CERN but at finer granularity varies

from 4 Gbit/sec to zero seen repeatedly. Not thought to be due to storage systems so the network people are keen to follow this up.

- KIT: 1) Both tape libraries are currently down. External technicians have been called but we cannot give a restoration time yet. 2) Our ATLAS SRM downtime finished at 9.20 am but SAM still thinks we are down although it is correct in GocDB.
- BNL: The slow CondorG job submission reported yesterday has been looked at by the developers who have found a workaround and will be updating their distributions. Job submission rates are now back to normal.
- RAL: 1) The URL for the Post Mortem of the problems with the migration of the Castor Oracle Databases at the RAL Tier1 can be found at: http://www.gridpp.ac.uk/wiki/RAL_Tier1_Incident_20100129 2) Had a glitch on the castoratlas DB this morning - quickly fixed. Also in a castor at risk just now changing their information provider. 2) Will have a major outage on 9 February changing site networking with two breaks of 15-20 minutes during a 90 minute period. FTS queues will be drained and batch will be paused. 3) There is an APEL (accounting DB) problem at the moment where the repository broke. Data is arriving but cannot be examined at the portal and SAM tests are failing. The repository is being reloaded from backup but the 100GB restore will take about 30 hours.
- IN2P3: 1) From yesterday's CMS reports of transfer problems from us to RALPP the problem seems to be solved but we cannot find tickets anywhere to allow us to investigate further. Pepe reported there was a Savannah report and he will send in the number. 2) Also from yesterday the LHCb issue of the extended downtime notice coming too late was because the person concerned did not realise that the GocDB had a built in notification delay of an hour to allow for late updates and the extension was entered just before it started rather than earlier. 3) We are having problems using YP (yellow pages) NIS from their worker nodes and do any other sites have any experience. Angela reported that KIT had migrated from YP to automounter and would check what was behind that.
- ASGC: 1) The power downtime on 9th Feb reported yesterday is now probably moved to 10 February. This will parallelise the systems to reduce the affects of scheduled safety checks and suchlike. 2) Repacking of tape space has freed another 40 TB of tape space making a total of 360 TB available to CMS.
- CERN FTS Pilot: The FTS pilot at CERN was upgraded yesterday to PATCH:59955 as requested by atlas yesterday.
- CERN SCAS pilot: A new SCAS authorisation service (a server replacing LCAS) is being started and needs to be tested in the experiment pilot job frameworks so please contact Gavin.Mccance@cernNOSPAMPLEASE.ch if you are interested.
- CERN CASTOR: 1) The network team has found a problem with a switch connecting eight CMSCAF servers and one T1TRANSFER server. To fix the problem the network switch will have to be rebooted causing the nine servers on CASTORCMS to be unavailable for a short period (~15 minutes). The proposal is to do this intervention tomorrow morning at 9am, this will cause some of the data on these pools to become unavailable for a short period of time. 2) A patch release for CASTOR 2.9-4 and also xrootd is available and is important for ATLAS being a prereq for SRM 2.9. The proposal is to migrate castort3 (shared with CMS but low activity) on Monday then castoratlas on Wednesday. The Oracle security patches would also be applied. They will do the other experiments later by negotiation.

AOB: Summary (and discussion) of this 1st GGUS post-change alarm test (G.Grein - from <http://savannah.cern.ch/support/?111475>)

- Next time we won't do all tests at the same time but split them into 3 slices: Asia/Pacific right after the release, European sites early afternoon (~12:00 UTC), US sites and Canada late afternoon (~ 15:30 UTC).

- We discovered some bugs for sites BNL, FNAL and NIKHEF which are already fixed. So the test was successful in this regard.

Jon Bakken queried this as nothing had been seen at FNAL. Gunter explained this had been a site naming issue internal to the GGUS system.

- The alarm for CERN-PROD (alice) is still open and seems not to be handled:
https://gus.fzk.de/ws/ticket_info.php?ticket=55244

- It looks like the alarm process at BNL is not clear to everybody. The alarm ticket was accepted quickly but the closing took about 6 hours: https://gus.fzk.de/ws/ticket_info.php?ticket=55263

Michael reported that BNL had neither received the alarm email nor the SMS that should go to his Blackberry though he knew OSG had got the ticket as he saw it having been closed. Gunter explained these three operations were done in parallel so he would check within GGUS in particular what addresses were used in this test. A follow-up test should be made. Miguel added that the CERN SMS gateway does not work for digitally signed messages.

(MariaDZ) I have a meeting clash but I updated
https://twiki.cern.ch/twiki/bin/view/EGEE/SA1_USAG#Periodic_ALARM_ticket_testing_ru to reflect the special measures on Tier1s' timezones.

Friday

Attendance: local(Miguel, Jamie, Maria, Edoardo, Nicolo, Dirk, Lola, Harry, Sveto, Timur, Jan, Roberto, Jean-Philippe Ale, Ueda, Andrea);remote(Xavier, Jon, Gonzalo, Onno, Rolf, Michael, John, Gang, Jeremy, Rob).

Experiments round table:

- ATLAS reports -
 1. No major issue to report
 2. RAL - We appreciate the announcement about the glitch
 3. CASTOR - We agree to the proposed upgrades on Monday (local analysis - CASTORT3, also used by CMS) and Wednesday (CASTORATLAS) next week. [Miguel - not on NS on CASTORATLAS. Part on Oracle security upgrade + CASTOR patch incl. schema change - 2 hours including xroot update and others.]

Rolf - we found that during last 24h ATLAS submitted as many jobs as all other experiments (LHC + non-LHC) together in last week. >50% have CPU time consumption <25s. Tons of pilot jobs? If so why?

Ale - will check. Production right now and transfers too at normal level. Please send some more detailed info.

Rolf - Yes, will open a ticket.

- CMS reports -
 - ◆ T0:
 1. Short unavailability of CASTORCMS disk servers on cmscaf and t1transfer for network switch intervention
 2. Some express stream jobs not starting on worker node or not reported as finished - GGUS #55295

3. Found FTS delegation problem but due to fact one of servers behind alias had not been patched. Fixed now.
- ◆ T1s:
 1. Backfill running at CNAF, IN2P3, PIC
 - ◆ PIC
 1. Service certificate used by central DataOps team failed to stage files as there was a mis-configuration at dcache level ([www.dcache.org #5419] Tape protection configuration reload). In the current release, one has to specify an special config entry ["*" "/cms(/.*)"] to catch all CMS roles. Patch applied at PIC and fix verified, ticket closed.
 - ◆ T2 highlights
 - ◇ MC Production completed in RAL T2 region
 - ◇ New generator-level events in production at CERN to replace datasets with duplicated events.
 - ◇ Some T2 tickets in full report
 - ALICE reports [↗](#) -
 - ◆ T0 site and T1 sites
 - ◇ Both CERN and T1 sites are currently performing a Pass4 reconstruction of raw data to test new updates included in AliRoot. ALICE reported a good behavior of services while performing this reconstruction
 - ◆ T2 sites - are continuing the MC production in parallel to the mentioned Pass4 reconstruction with no incidents to report for these sites
 - LHCb reports - Several MC production launched yesterday evening and now running at a low rate of few thousand jobs reached a peak of ~9k concurrently running MC simulation jobs last night. In the picture (see full report) the evidence of this activity over the last 24h. as reported by the SSB.
 - ◆ T0 sites issues:
 - ◇ SRM is unusable: LHCb open a GGUS ticket because noticed that SAM jobs failing since this morning ~3:00 am with all SRM requests timing out. Jan was sending at the same time a mail reporting about the funny state the SRM endpoint was. There are a couple of possible reasons behind but LHCb do not believe that is putting in the system a load much larger than in the past (despite the increased activities in the last 24 hours):
 1. high load on 'lhcbdata' by user "lhcbprod" (2.4k outstanding requests, these seem to write rather slowly (not at wire speed) to the handful of servers that have space left) [Jan - migrator problem? Restarted yesterday morning at 11:00. At night many files for migration]
 2. SRM-2.8 bug that makes SRM a single point of failure over all pools <https://savannah.cern.ch/bugs/?45082> [↗](#) (fixed in 2.9, but that needs further validation) - otherwise this load would only fails access to "lhcbdata", but now it affects other pools as well. [Jan - known feature / bug. Will put on PPS SRM for ATLAS - can be done also for LHCb]
 - ◆ T1 sites issues:
 - ◇ IN2p3: SRM endpoint became unresponsive: both SAM tests and normal activity from our data manager were failing with the error. The suspicious is that some CA certificate is not properly updated on the remote SRM. In this case the CERN CA.
 - ◇ RAL: a Oracle glitch reported by our contact person immediately fixed by RAL but that might have caused some jobs failing/stalling due to a temporarily service unavailability.

Sites / Services round table:

- CERN - intervention on network switch. Was this scheduled? A - yes, was agreed in advance.

- CERN DB - data loss to BNL still under investigation.
- CERN FTS - on Monday we will release for a few sites FTS 2.2.3 which contains fix for proxy failure. But doesn't contain all FTS modules. Need to install FTS 2.2.2 and overwrite some modules by FTS 2.2.3. On Monday will probably tell sites with FTS 2.2.2 to go to new version and wait for T1SCM on Thursday for other sites. Michael - do sites know about this? JPB - will update on Monday if we are still happy on Monday: BNL, FNAL and KIT are sites currently running FTS 2.2.2 that could update if the green light is given on Monday.
- KIT - ntr
- FNAL - ntr
- PIC - ntr
- NL-T1 - ntr
- IN2P3 - ntr
- BNL - ntr
- RAL - DB glitch: an node that was part of a RAC rebooted. Was not in use - spare - but when rebooted Oracle decided to redistribute load - caused 10' outage. DB people looking into it - should not have happened. Lost all CASTOR activity for 10' NS and ATLAS + LHCb. A couple of rogue WNS. Rebuilt but had misconfig. Stalling jobs - removed WNs from production when noticed (after ~1h).
- ASGC - ntr
- GridPP - ntr
- OSG - 1 ticket that has been hanging around for some time - 54260 - against BNL. Question from user back in December. Still marked very urgent but answer by John Hover. If user thinks still urgent can reopen.

AOB:

-- JamieShiers - 29-Jan-2010

This topic: LCG > WLCGDailyMeetingsWeek100201

Topic revision: r21 - 2010-05-27 - SteveTraylen



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