

Table of Contents

Week of 100118.....	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.....	6
Friday.....	8

Week of 100118

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(Jamie, Maria, Harry(chair), Nicolo, Andrew, Jean-Philippe, David, Alessandro, Jaroslava, Gavin, Ignacio, Ricardo, Simone, Patricia, Eva, Roberto, Julia, Dirk);remote(Gonzalo(PIC), Angela(KIT), Michael(BNL), Onno(NL-T1), Gang+Jason(ASGC), Pepe(PIC+CMS) ,Daniele(CMS), Mark(IN2P3), Gareth(RAL)).

Experiments round table:

- ATLAS reports - 1) Problems with NDGF FTS service (out of space in /var). Quick intervention from T1 to alleviate the immediate problem but cause still under investigation. 2) Tier-2s of Napoli, Lisbon

and QMC-London down over the weekend.

- CMS reports - 1) Pending jobs revealed that some hosts on CMS T0 apparently are never running CMS jobs, already seen in December but at that time the unused hosts were in maintenance. About 500 job slots so being followed up. 2) New round of backfill jobs submitted to test CREAM CEs, issues at CCIN2P3 and CNAF reported last week were fixed. 3) Deployed debugging library at FNAL to track rare error in opening files on dCache. 4) SAM CE test failures at ASGC this morning after a power surge.
- ALICE reports [↗](#) - 1) Four pending requests for myproxy registration of new Tier-2 VO-boxes have now been completed. 2) Production ongoing at all T1 sites with no incidents to report (scheduled downtime at SARA). 3) High overload of both ALICE VOBOXES at IN2P3 reported on Friday solved on Saturday morning. The CREAM-VOBOX entered production at that moment. No interruptions of the services observed. Good performance seen.
- LHCb reports - 1) There are 4 new productions requested and all are about full reconstruction of COLLISION09 data. Only users during the week end and now less than 1000 jobs in the system. 2) Issue with SRM unit tests failing at some sites. Working for a solution. 3) Waiting for news for a second box at CERN behind LFC-RO. 4) Intervention to migrate CASTOR to TSM at INFN-T1 (6TB in total, ~2K files). They are all ready to proceed. Storm and CASTOR endpoints banned at CNAF to proceed with. 5) Angela reported from KIT on the CREAM-CE mapping problem seen by LHCb last week. There were wildcard entries (*) in the mapping configuration that CREAM cannot work with. Now removed - LHCb will test again.

Sites / Services round table:

- PIC: Reminder that there will be a cooling intervention tomorrow with batch capacity reduced down to about 400 slots.
- KIT: A sudden increase of ATLAS transfers to Wuppertal filled the FTS disks to over 50% with logs which had to be cleaned - a cron job has since been added but this activity was not expected. Simone reported this was FTS traffic to a group file space where the issue was not bandwidth but the number of files, about 70K, and he suggested to retune the FTS settings. Not clear this would help since the issue was log files filling space.
- BNL: Had a network problem Sunday when a 10 Gbit ethernet line card on a core router failed and had to be replaced. Some jobs failed and were put by PanDA into the "holding" state. They will fail eventually later with "lost heartbeat".
- NL-T1: Nikhef is draining batch queues for a torque and bdi upgrade tomorrow afternoon. At SARA all nodes have been rebooted following the scheduled kernel upgrades. Roberto queried SAM test failures at SARA for LHCb today and Onno reported that SRM had not been restarted properly but is now fixed
- ASGC: They suffered a power surge about 07.25 UTC which took down most core services. Most were recovered within 30 minutes but Oracle critical services took 2 hours. Job scheduling is back to normal but the CASTOR service and SRM are not yet restored - expected to take another hour.
- RAL: Last Friday's FTS problem was resolved by the end of the afternoon though there was another problem for 10 minutes this morning when a database node behind the service was rebooting. There will be at-risks scheduled this week: CASTOR SRM updates tomorrow and Wednesday and a memory upgrade to Oracle RAC nodes behind CASTOR.
- IN2P3: Currently have a problem of worker nodes being stressed by jobs and ending up being

rebooted. The suspicion is on ALICE jobs which have hence been blocked in batch for now. There were CREAM-CE problems at the weekend which may be connected with this. There was also an SRM failure at the weekend where it was down for 17 hours leading to SAM test failures. A script is being prepared to automate recovery.

- CERN Databases: The latest Oracle security patches are being deployed on the integration services.
- CERN dashboards: The poor performance of the CMS dashboard has finally been resolved as due to memory caching. The memory has been doubled and performance has improved threefold.
- CERN networking: A 2.5 Gbit/sec traffic from Uni Geneva overloaded the firewall and had to be stopped. The suggestion is for them to register to use the firewall bypass. Simone reported this as due to a single user downloading a dataset that was only in Geneva. It should have been done via subscription and letting FTS move it rather than by direct copy but such actions in individual Tier-2 sites are unpredictable and might happen at any time. Eduardo pointed out that Geneva is special in having a 10 Gbit link but it was agreed this issue needs following up.
- CERN FTS: The issue with corrupted FTS delegated credentials is now understood thanks to FNAL. It turned the longstanding race condition bug was erroneously restored in a build of FTS 2.2 by including an old dependency. A new version, including also a fix for the bug of agents crashing, will be released as soon as possible.
- CERN CASTOR: The castort3 (CMS and ATLAS analysis stager) was upgraded this morning during a 4 hour scheduled downtime.

AOB: (MariaD) On GGUS-OSG issues: The following is taken from last Monday's 2010-01-10 wlcg notes: 'Could OSG please update https://gus.fzk.de/ws/ticket_info.php?ticket=54538 (urgent) Michael clarifies: Harvard and Boston form together one T2 center. MariaD will discuss ticket routing for this case offline with Kyle and Michael. ' Today this ticket is still urgent and untouched. Comment by Guenter: if there is still no corresponding OSG ticket, someone should do the following:

- assign ticket to GGUS
- assign ticket to OSG again.

Michael (BNL) reported that the problem had long been resolved and what remained was to solve the consistency issue of Harvard and Boston being essentially a single Tier-2. Should be solved in the next few days.

The ALARM monthly text conclusion is 2010-02-03 pm will be the 1st testing date, right after the Jan GGUS release. Doc updated https://twiki.cern.ch/twiki/bin/view/EGEE/SA1_USAG#Periodic_ALARM_ticket_testing_ru Details: <https://savannah.cern.ch/support/?111475#comment13>

Tuesday:

Attendance: local(Jaroslava, Harry(chair), Nicolo, AndreaV, Roberto, Jean-Philippe, Eva, Patricia, Lola, Alessandro, AndreaS, Jamie, Julia, TimurS);remote(Gonzalo(PIC), Angela(KIT), Jens(NDGF), John(RAL), Ronald(NL-T1), Pepe(PIC&CMS), Jeremy(GridPP), Gang, Jason).

Experiments round table:

- ATLAS reports - 1) SARA/NIKHEF back from downtime yesterday before noon. FTS did not come back but was fixed a few hours afterward. NL cloud back online from this morning. 2) Wuppertal SE was down in the morning so they restarted their SRM door. However the FTS channel had meanwhile

been set offline by FZK to reduce errors and had to be reopened to restart data transfers. 3) Jens queried the reason for the large amount of transfers seen at NDGF over the weekend and Alessandro reported it as user activity that must be expected from now on.

- CMS reports - Daily report: 1) IN2P3 suffering in CREAM CE tests from a dcache issue in stageout since last week. According to local contact, the issue was identified, but the final fix is not available, and a workaround will be put in place soon. 2) KIT experiencing 'Error 73 - Failed writing to read-only file system.' in CREAM CE tests. 3) Dashboard reporting failures in CREAM CE tests at CNAF, no detailed logs available in CMS submission framework yet. Also problems staging in from tape delaying data export from CNAF. 4) FNAL error reported yesterday in opening files on dCache is possibly related to too many file descriptors open. 5) SAM CE test failures reported yesterday at ASGC are fixed but there are now SRM test failures. 6) Julia reported on mismatches between CMS running jobs and the number reported by the dashboard as due to many jobs staying in the WMS-LB system for hours. Weekly Planning: 1) Tier-0 running workflow and release tests. 2) Tier-1 running backfill testing jobs at all T1s (except ASGC). 3) Tier-2 to finish the two running Madgraph workflows. Might receive new requests this week. 4) Following-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. 5) Following-up CREAM CE tests and failures observed at several Tier-1s. 6) 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. 7) CMS first mid week global run of 2010 (MWGR) to occur end of next week (27/28-Jan).
- ALICE reports [↗](#) - The ALICE contact person at IN2P3 has been contacted in order to collect more information about possible ALICE jobs blocking the local WNs with an excess of i/o operations. From ALICE side, the experiment did not change the production type over the weekend and the only additional element is indeed the submission through CREAM. IN2P3 will set up a specially monitored subcluster and redirect ALICE jobs there.
- LHCb reports - 1) Ongoing activities in certifying the new DIRAC production system (based on SVN) over the new h/w delivered for central boxes. 2) The CERN volhcb13 server got its dirac partition full and then the job logging info database corrupted causing some user problems in recovering outputs. 3) GRIDKA confirmed that now CREAMCE are mapping correctly sgm users. 4) at CNAF Storm is being upgraded to a recent version for supporting TSM. 5) SRM SAM jobs were failing at RAL. This seems related to the new code of the unit test from LHCb which is now debugged. Currently rolled back to the old-stable code for the critical unit test. 6) There is another failed diskserver at RAL in the lhcbDst space token. John reported a disk server containing 3200 files had temporarily been taken out of production yesterday for a memory check. He reported there was another LHCb disk server giving fsprobe errors that was being worked on but Roberto confirmed this was an older problem. 7) Jean-Phillipe had followed up on the LHCb request for another LFC server behind CERN Read-only instance and asked them to submit a ticket - earliest delivery will be next week. Remedy ticket was open on Monday (CT654872)

Sites / Services round table:

- NDGF: Close to the end of an FTS update. Also added a disk to cope with future numbers of FTS log files following the large number (1 million logs from 600000 transfers) seen over the weekend.
- NL-T1: Restarting after the SARA maintenance led to two problems - the reboot of the ATLAS LFC got stuck and the FTS version was accidentally upgraded to version 2.2 but has now been returned to 2.1. At NIKHEF the bdii migration to a 64-bit platform is completed as are the kernel upgrades and the torque upgrade is ongoing.
- GridPP: The London Royal Holloway Tier-2 (ATLAS and CMS) will be in extended downtime for the next 4 weeks.

Tuesday:

- BNL (email report): Unable to attend due to overlapping meeting but no particular issues.
- ASGC: Following the power surge event yesterday we found the srm transfer efficiency reduced to 60-70% only. It affects atlas production transfers and cms also observe the instability of srm transfers today. Since yesterday, we are able to fix: * direct file i/o able to complete shortly. * scheduling problem also have been fixed after restarting the rm master, and also job manager after force killing the ipc processes. * confirm the full func of stager and also rh service. * we found that the SRM has sent the request and it was received by the rhserver but never picked up by the stager. * have recompute the stats yesterday. * now we reduce the stage timeout to 90s, and remain the same at 180s for file transfer for srm timeout.

AOB:

Wednesday

Attendance: local(Jaroslava, Harry(chair), Nicolo, Jamie, MariaG, Jean-Philippe, TimurS, Simone, Dirk, Steve, Denise, Eduardo, Patricia, Roberto, MariaD, Miguel);remote(Angela(KIT), Ginzalo(PIC), Ron(NL-T1), Gang+Jason(ASGC), John(RAL), Rolf(IN2P3), Jens(NDGF)).

Experiments round table:

- ATLAS reports - Nothing to report today.
- CMS reports - 1) CERN CASTORCMS DEFAULT degraded last night 23:00-01:00. 2) Investigating failures in CREAM CE tests at IN2P3. 3) KIT '426 Transfer aborted' error in FNAL-->IN2P3 4) ASGC - SRM issue fixed. 5) Software deployment on SLC5 starting on T2_UK_London_IC. 6) still 3 T2s with open tickets. 7) T2_EE_Estonia lost NFS software area, also reflected in SAM tests
- ALICE reports [↗](#) - 1) Change of general production cycle so an unstable job profile should be expected for today. 2) Ticket sent to PX support from Birmingham by mistake CT655327, please skip it. 3) Ticket sent to VOBOX support: CT655195. ssh access to the ALICE VOBOXES is closed. The issue might come from a change in the default behavior of the ncm-useraccess component which was announced in December and went into production with the last scheduled upgrades (last week?). Operational use of the machines is ensured (via gsissh), any sudo operation is however denied. Following with the experts. 4) Following the CCIN2P3 issue reported since Monday the site admin locked every aligridd job. At about 16:00 the jobs were unlocked in a controlled way by lowering BQS limit during the observation period (until today), this means setting the limit to 300 slots. Once the slots were opened, ALICE closed the submission to the local CREAM-CE system just to ensure a single submission procedure, using therefore the WMS submission mode only. We have asked the site admin if the problems observed on Monday are still visible and if we can open again the submission to CREAM. Rolf later reported that the analysis of the problematic ALICE jobs is ongoing but it could just be a coincidence as regards ALICE.
- LHCb reports - 1) LFC 3D replication out of CERN yesterday had some problem(trapped by our SAM and SLS and confirmed by Eva who received an alarm). At 16:47 we got an alarm from our monitoring system: the capture process was delayed for 90 minutes. There was a problem receiving the archive log files from the source database on the downstream system and a GGUS team ticket was submitted. 2) LHCb provisionally agree on the intervention for the CERN CASTORLHCb upgrade to 2.1.9 the 26th, 3) CNAF: Storm successfully upgraded (1.4--> 1.5) and SAM confirmed it is back to life. Data migrated to TSM. Now changing the configuration service and adding these replica in LFC. The TxD1 endpoints have been reintegrated.

Sites / Services round table:

- PIC: 1) ATLAS have submitted a ticket on transfers from NDGF to PIC failing with timeout errors. We have asked NDGF if they have changed something as PIC have not but this problem may have been there since many days and is perhaps related to the previous trial of using jumbo frames. 2) Since several days LHCb jobs have been running at very low cpu efficiency seemingly waiting for something. Roberto thought this must be individual user jobs and will check.
- NL-T1: 1) NIKHEF have increased the MTU on their router to the LHCOPN from 1518 to 9018. 2) NIKHEF also has an issue with their MAUI client which cannot talk to the MAUI server after an upgrade - being worked on. 3) At SARA one pool node is down when both disks of a RAID-1 array failed at the same time but with no data loss. Being reinstalled and should come back later today.
- ASGC (email report): The CASTOR problems root cause arise from bad service assignment at backend db cluster. All the services should split into three nodes from the castor db cluster but after the power cycle, all the services fall back into same node thus we have observe extreme load situation at this particular node. Have force altering the distribution and have the other services taking the free node. Furthermore, after the power cycle, seems non-smp kernel is favored after system reboot and this is causing limited performance from current dual core instances of Oracle cluster. Have gradually reboot all 5 nodes from two cluster, this actually improve a lot the overall situation. I try simple service rearrangement 20min ago, and it actually help improving the situation, up to 83%, and I am restarting the srm and stg service right after the rearrangement of the other services. Am expecting the transfers will be able to stabilize further after the change,

Release report: deployment status wiki page

AOB: 1) (MariaD) USAG tomorrow Thursday 2010-01-21 at 9:30am in room 28-R-014 and by EVO. Two major developments to be discussed - see Agenda <http://indico.cern.ch/conferenceDisplay.py?confId=81363>. 2) (MariaG) The first of a new series of WLCG weekly service coordination meetings will be held tomorrow from 16.00. The mandate and agenda are at <http://indico.cern.ch/categoryDisplay.py?categId=2726>

Thursday

Attendance: local(Jamie, Maria, Gavin, Ricardo, Lola, Nicolo, Jean-Philippe, Timur, Julia, Przemyslaw, Jaroslava, Roberto);remote(Gareth, Jens, Angela, Gonzalo, Ronald, Jason, Gang, Jason, Rolf).

Experiments round table:

- ATLAS reports - ASGC not getting FTs, FTS looks stucked, elog:8785, GGUS:54837; NIKHEF-ELPROD: put back into DDM on atlddm29, now is whitelisted on every SS, elog:8797; TORONTO-LCG2: site suspended by CA ROC, stay blacklisted in DDM, elog:8782, DDM-savannah:61670; IN2P3-LPSC_DATADISK put back to FT, elog:8792
- CMS reports - Lost interactive access to WNs; T1s: PIC CMS filling completely job slots (stress tests); CNAF - all CMS resources now SLC5; ASGC: some instabilities with SAM tests but now green again. Transfer errors FNAL-IN2P3 reported yesterday were due to a few 35GB files. Temporary solution until FTS2.2 is to increase timeout, transfer large files and then reduce. T2s: Another DPM site (T2_FR_IPHC) hit by SLC5 CMSSW-DPM compatibility issue - reminder: currently the SLC4 compatibility libraries need to be installed on the WNs of CMS DPM sites even if only native SLC5 software builds are used, documented here: <https://hypernews.cern.ch/HyperNews/CMS/get/sc4/2085/1/2/2.html>; T2_EE_Estonia NFS software area recreated, need to recover CMS site-local-config, reflected in SAM tests; T2_BR_UERJ - SRM not visible and CMS services down since a few days; MC production in RAL T2 region progressing. IN2P3: worried about story of large files as this occurred frequently in the past - had to have larger timeouts. Hope for solution viable for our FTS people who are frequently asked to redefine timeouts.

Would like to ask CMS to take into account the fact that FTS channels are in principle shared (not with FNAL...) If there is a real problem have to wait for very long due to timeout length. This could create real problems when LHC producing data in full stream. Gav: FTS 2.2 contains capability for more elastic timeouts based on filesize - need to understand how to tune this. Nicolo: protections have been put in place in T0 to avoid regular production of files > 10GB - these cause problems on some WNS - maybe not enough local disk space e.g. on multi-core WNs. In future should not be so bad.. IN2P3 retransferring data lost from last year so many files > safe limit. Will test FTS 2.2 timeouts when deployed. Worth looking at transfer speeds - can this be improved without changing timeouts? Improvement here could also allow a larger max filesize without extending too much the timeouts. Maria - should follow this at T1 service coordination meeting which kicks-off today.

- ALICE reports [↗](#) (Reported by Patricia before the meeting) - Problem with the access to the ALICE VOBOXES at the T0 solved. Once the changes in the VOBOXES templates are included, it takes a while until its full distribution to all nodes. Now both the sudo privileges and the access to the machine are working fine. Ricardo has been informed and the ticket can be closed. Clermont and Strasbourg mentioned today bad results visible in the MonaLisa page, regarding the publication of the user proxy status. The issue is coming from the script responsible of the testing of this proxy. Differences in the libs used by this script (Alien environment) and the libs of the VOBOX service itself created such bad (and fake) results. Solved in both sites, the propagation to the rest of the sites will be discussed today during the ALICE TF Meeting
- LHCb reports - Reprocessing of 2009 450 GeV data launched. LOHLT MB MC09 Stripping finished and a summary available here : CNAF and GRIDKA (as reported last week) the most problematic sites. The bb and cc inclusive stripping in the system now.No T0 issues; T1 issues: CNAF: registering data of the new T1Dx endpoint in LFC.; PIC: The low efficiency jobs observed yesterday (as suspected) were user jobs: these were about jobs whose output sandbox upload to CASTOR RAL was hanging. DIRAC has in place any possible timeout. Jobs stack is not longer available because finally killed by the LRMS no further investigation are possible.; RAL: similarly another user job seems to have consumed 49 seconds over a wall clock time of 55 hours,.

Sites / Services round table:

- KIT: A few problems with ATLAS user jobs - asked by local ATLAS admins to block one user but not problem with jobs but framework - too many dcap connections. Seems solved now.
- PIC: Issue reported yesterday NDGF transfers getting stuck still open - still unclear. 1/2 transfers failing with timeout. Probably not jumbo frames. Timeout uploading sandbox into CASTOR at RAL - is this standard procedure? Robe - depends on workflow of user job. Can happen.
- NL-T1: ntr
- IN2P3: some news about ALICE - incidents which crashed a lot of WNs. Found that 3 users were responsible for this but apparently was combination of these users. One was mapped login for ALICE jobs. ALICE can now submit jobs as usual. Normally should not result in any problems by themselves. Apparently a combination with other jobs from QCD analysis jobs- very high I/O rate - hope to separate both. Issue ongoing...
- ASGC: already fully recovered since 20:00 yesterday - all SRM tests etc ok now. New CE published with 1500 cores as SL5 cores. Nico - repacking news? A: no
- RAL: main thing: short outage of FTS due to h/w fault on node where agents run. FTS drained and agents moved. Advance warning - some significant outage next week: details in GOCDDB.
- NDGF: ntr

- CERN: DBs: ALICE Online DB - upgrade of firmware on storage - should be finished in 1-2 hours. Switch between standby and primary planned next Thursday during UPS intervention.
- CERN: new kernel vulnerability for RHEL and SLC5 kernels. Will be rebooting m/cs at short notice - lplus and batch capacity. Degradation in available capacity during this. Will start with lplus.

AOB:

Friday

Attendance: local(Jaroslava, Hary(chair), Lola, Roberto, Jamie, Ricardo, Timur, Eva, JanI, Nicolò);remote(Jens(NDGF), Xavier(KIT), Gonzalo(PIC), Tristan(NL-T1), Rolf, John(RAL), Michael(BNL), Jeremy(GridPP)).

Experiments round table:

- ATLAS reports - 1) Ran out of disk space on Nikhef DATADISK but basically an ATLAS issue that should be solved very soon. 2) There was a monitoring issue with CERN SLS where around 13.00 shibboleth authorisation stopped working and the daemon had to be restarted by SLS support. Timeouts were seen around 14.00 on the Atlas distributed computing web pages.
- CMS reports - 1) T0 operators lost interactive access on cmst0 worker nodes (needed for debugging troublesome workflows), investigation ongoing. 2) New CERN SLC5 VOBOXes vocms02,vocms03 available for PhEDEx, requested registration in myproxy. 3) Large files in FNAL-->IN2P3 transferred after timeout increase on FTS channel, follow up discussion to be held. 4) Running pre-production testing at IN2P3 for a new round of MinimumBias events reconstruction. 5) Some files exported from ASGC arrived with with an incorrect checksum which might have been corruption at the ASGC source. Gang remarked that he was doing some checking and noticed that ASGC people were not included in the Savanna notification of this problem - Nicolò will follow this up. 6) Started deployment of SLC5 software releases at ASGC.
- ALICE reports [↗](#) - 1) Change of production cycle still ongoing. Unstable job profile should be expected for today. 2) During yesterday's ALICE TF Meeting the final list of sites which have been blacklisted due to a non upgrade of resources to SL5 has been presented. This activity has been completed 3) A CERN CASTORALICE intervention has been announced by the CASTOR team for the next week (27th Wed). This intervention will upgrade CASTOR software to version to 2.1.9-3-1. This upgrade requires a service downtime. It will start at 9:30am Geneva time. The intervention window is 4 hours. Both ALICE Offline and Online responsible have agreed with the intervention and the time window.
- LHCb reports - 1) CNAF is still in the middle of a migration (LFC catalog to be updated according). 2) IN2P3 preferred to postpone the test of the new gsidcap client to next week.

Sites / Services round table:

- KIT: Announced a downtime of the ATLAS dcache from 1 to 5 February for migration from PNFS to Chimera so no access to ATLAS data for that week.
- PIC: The intermittent timeout failures in transferring Atlas data from NDGF to PIC are still ongoing. There were also issues from NDGF to FZK so the network experts are now looking for correlations.
- IN2P3: Announcing a downtime of the batch system on the 25th and 26 Jan to replace database servers. Then there will be no tape access on the 26th (writes will accumulate) and also an at-risk for a maintenance on the Oracle databases (LFC etc.).

WLCGDailyMeetingsWeek100118 < LCG < TWiki

- RAL: Ran out of ATLASCRATCHDISK space following many transfers. Two new servers have been added and monitoring has been improved. There will be a GOCDDB downtime next week.
- CERN Databases: An upgrade of the ALICE standby database has been postponed by ALICE due to a power cut at their pit. If they have time they will do it today.
- CERN Linux: Urgent kernel security upgrades were applied yesterday and batch nodes are being drained and rebooted - mostly already done at the Tier-0. This will likely be required WLCG wide.
- CERN SRM: There was instability on SRMATLAS overnight probably linked to high activity on the scratch disk at CERN putting too much load on the stager. This lasted about 90 minutes. SAM tests also reported a short outage on SRMPUBLIC during the morning, but this was a false positive (test ran without credentials).

AOB: Draft notes from yesterdays new format WLCG Service Control Meeting have been linked to the agenda at <http://indico.cern.ch/categoryDisplay.py?categId=2726>. The next one will be in 3 weeks on 11 February (normally 2 weeks but 4th is blocked off by other activities).

This topic: LCG > WLCGDailyMeetingsWeek100118

Topic revision: r17 - 2010-01-22 - JanIven



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