

# Table of Contents

<b>Week of 100215.....</b>	<b>1</b>
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.....	7
Friday.....	9

# Week of 100215

## LHC Operations

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

## WLCG Service Incidents, Interventions and Availability

VO Summaries of Site Availability				SIRs & Broadcasts	
<a href="#">ALICE</a>	<a href="#">ATLAS</a>	<a href="#">CMS</a>	<a href="#">LHCb</a>	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](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](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:

<a href="#">STEP09</a>	<a href="#">ATLAS</a>	<a href="#">ATLAS logbook</a>	<a href="#">CMS WLCG Blogs</a>
------------------------	-----------------------	-------------------------------	--------------------------------

## Monday:

Attendance: local(Harry, Gavin,Lola, Simone, Eva, JPB, Timor, Miguel, Ale, Jan, MariaD, Roberto, Dirk);remote(Jon/FNAL, Michael/BNL, Nicolo/CMS, Xavier/KIT,Gang/ASGC, sip, John/RAL, Rob/OSG, Vera/NDGF ).

Experiments round table:

- ATLAS reports - (Simone) reprocessing load created issues with several sites - please refer to ATLAS twiki for more detailed report. ASGC: FTS issues from Fri, fixed Sat. Need to upgrade from old FTS 2.0. PIC: disk pool crash on Fri caused 20% failures for exports. Checksum issues on transfer to BNL.

BNL: From Sat Oracle maximum connections reached: triggered LFC restart. Between Sat and Sun high load on PNFS (100k namespace ops/h - ~30 Hz - close to system limits). New reprocessing pattern results also in increased SRM load. RAL: low transfer speed to UK cloud and transfer failures due to offline CASTOR pool. TRIUMF: Sat FTS unreachable - fixed on Sun. DB outage created overload on failover DB - retuning DB memory. NDGF: slwo transfers of large files - Vera: Are Jumbo frames required? To be discussed eg at OPN meeting next week. Will contact OPN responsible and notify site. SARA: SRM issues - still open. CERN: client hung on xroot connections (Sun evening): Thread exhaustion fixed by xroot server restart but reasons still being investigated.

- CMS reports - (Nicolo) - T0: CASTORCMS low SLS availability on Saturday, maybe stager network issues, GGUS #55521. Miguel: CASTOR @ CERN experiencing intermittent timeouts - (also ATLAS) - investigating with network experts. T0 manager crash on Sunday for issues in connecting to DB, due to failures in CMSR cluster, recovered at 15:20 Remedy #661430. SLC5 gridftp server needed to test new CRABSERVER, ticket with IT-PES to provide missing rpms Remedy #CT0000000659265. T1s: Increased rate of backfill tests of CREAM CEs at European T1s - failures in loading CMSSW env at IN2P3, was fixed. IN2P3: Transfer failures from T0 and FNAL for proxy issue, Savannah #112741. FNAL: Batch farm issues solved at FNAL, skims completed. RAL: Skim jobs without LazyDownload causing saturation of storage-->wn network, killed. Some 155 files lost due to defective tapes. CNAF: Timeouts in import of MC from Vienna. ASGC: CASTOR files with invalid diskcopy restaged from tape. T2\_DE\_DESY reported incompatibility between gsidcap libraries in latest gLite release and CMSSW, switched back to dcap, CMS offline investigating update to libraries shipped with CMSSW
- ALICE reports - (Lola) T0: Good behavior of the services during the weekend. CNAF: The latest CREAM-CE service announced last week by the site admin has been put in production. Remarkable stability during the whole weekend achieving peaks over 1000 concurrent jobs with a single CREAM-CE.
- LHCb reports - Roberto - IVery low level user activity. Tomorrow at 17 DIRAC will be turned off and on Wednesday the major s/w and h/w upgrade. Saw @ CERN root protocol timing out - transient. Miguel: possibly similar to CMS/ATLAS problems. LHCb will sent failure logs.

#### Sites / Services round table:

- Jon/FNAL - ntr. Jean-Philippe: any plans to migrate to new FTS version - Jon: yes but need to do tests beforehand.
- Michael/BNL: planning to move to new FTS on Tue next week, but need instructions. JPB: will be provided. Sat: large number of connection between LFC and Oracle. - connection last for too long and therefore pile up. Now stopped restarting LFC above some connection number threshold. Stopped restarting but need explanation for large number of ongoing connections. JPB: could we get LFC logs and number of LFC threads configured? Michael: will be provided. dCache pnfs issues tracked to config problem - now resolved. Simone: after second problem in SRM -this morning we stopped deletion - should we restart? Michael: yes, please restart.
- Xavier/KIT: downtime planned for afternoon of Thu (4h): OS update on router and file servers. Only dCache will be down. JPB: FTS upgrade to latest version? Xavier: not this week. Is there is no official release? Yes, there is a release. KIT will follow up on FTS plans.
- Gang/ASGC: ntr
- John/RAL: problems for ATLAS transfer due to LSF problems - fixed now. CMS : 2 tapes unreadable result in loss of 155 files - post mortem will be produced. Planned at risk for tomorrow (DB memory upgrade for CASTOR DBs) may move to Wed - will confirm
- Vera/NDGF: ntr
- Eva/CERN: Node failures on PDBR (1 out of 4) due to file system corruption (node needs reinstall). Kernel issues on one node of CMSR. Fri afternoon - human mistake on switch config caused loss of storage connection - all integration DBs rebooted. Investigating replication delay of 2h between

ATLAS online / offline DB.

- Miguel/CERN: CMS CASTOR upgrade to 2.1.9-4 as scheduled (DB security patch, SRM upgrade). Same Wed on CASTOR public - CERN site needs to declare downtime for Wed intervention, but other CASTOR services will stay available.

AOB: (MariaDZ) On RobQ's question from last Friday (also sent to him by email): The answer to who retires the ticket (I hope by retire you mean solve) is: The Support Unit to which it is assigned. This is valid for all tickets. In case of ALARM tickets, this is the ROC to which the Tier1 belongs. In case of OSG tickets, this is the OSG GOC. GGUS will launch an alarm to FNAL and BNL today as of 4pm CET to test the additional feature implemented for Rob Quick. Any objections, please say so now. Proposed alarm test has been confirmed with Jon and Michael.

## Tuesday:

Attendance: local(Jamie, Katarzyna, Timur, Lola, Jan, Ale, Simone, Andrea, Roberto);remote(Nicolo, Jon, Gonzalo, Angela, Rob, Daniele, Michael, Gang, Tristan, John, Jeremy, Rolf, Maria, Vera).

Experiments round table:

- ATLAS reports -
  1. SARA:
    - ◇ SRM timeouts due to server overload. Unscheduled downtime marked this morning. Now seems fixed, please confirm.
  2. CERN:
    - ◇ In downtime for the upgrade of CASTOR-public. But srm://srm-atlas.cern.ch:8443 is also marked in downtime. This was explained yesterday, but was not completely clear to me. The downtime would have side effect if ATLAS would use downtime information in an automatic way, which will be stating doing soon. Jan - should be able to get away from this when we have rolled out new site tests. Existing tests use CASTORPUBLIC - new ones will use experiment. Ale - not a matter of tests but of publication of downtime. Cannot be done today for different end-points. Propose to take offline Kasia - PVSS replication delay of 2h reported and now solved. Due to problems with huge tables.
- CMS reports -
  - ◆ T0
    1. CASTORCMS back after upgrade, no issues seen so far
    2. Short drop of CASTORCMS total space in SLS - SLS glitch or actual issue? Remedy #CT0000000661844
    3. SLC5 gridftp server needed to test new CRABSERVER, ticket with IT-PES asking for CDB template Remedy #CT0000000659265
  - ◆ T1s
    1. MC reprocessing at PIC, IN2P3
    2. Skimming tests at RAL, FNAL, IN2P3
    3. Reduced rate of backfill tests of CREAM CEs at European T1s.
  - ◆ CNAF
    1. More timeouts in import of MC from Vienna. (Ticket will be updated)
  - ◆ T2 highlights
    - ◇ MC ongoing in IN2P3, PIC, FNAL T2 regions
      1. Progress in ProdAgent plugin for submission to NorduGrid T2 (Helsinki)
    - ◇ Ongoing SAM test failures at T2\_IN\_TIFR.
  - ◆ CMS Weekly Scope operations plan
    - ◇ [Data Ops]
      - T0: support global run, process and distribute data.

- T1: run backfill testing, waiting for other real requests
- T2: complete re-production of LHE requests
- ◇ [Facilities Ops]
  - Follow-up on Tape Families setup in ASGC to bring the site to Operations before the run. Still waiting site's green light to start with backfill tests.
  - A clash using gsidcap and CMSSW latest versions seen. CMS itself has not imposed any specific requirements on the sites to use fully authenticated protocols for (read-only) data access to CMS data, thus in principle CMS is also not requiring anything more than simple dcap access. We are evaluating to ask all CMS sites they drop back to normal dcap. More news soon.
- ALICE reports [↗](#) - GENERAL INFORMATION: All MC cycles have finished successfully and for the moment there are no requests of new cycles inside MonaLisa. However, the production is ongoing with user analysis jobs. At this moment, there are 2 analysis trains. In addition, ALICE is performing the reconstruction of raw data (pass1 and pass2 reconstruction of cosmics) at T0 and T1 sites
  - ◆ T0 site - CREAM-CE system has been tested again this morning. No issues have been found at CERN
  - ◆ T1 sites
    - ◇ Performing the pass2 reconstruction at FZK and RAL at this moment. No issues to report in terms of services at these sites
    - ◇ AliEn services restarted this morning at CNAF to restart the production at this site
- LHCb reports - All activities of production team are focused on preparing the major migration to DIRACv5r0. This implies testing new (already certificated) software in its new configuration, running over new machines. DIRAC will be switched off today at 17:00 (if this pre-production like activity does not show up any major problem)
  - ◆ T0 sites issues: Open a low priority ticket to track down the timeout issue accessing data via root at CERN with some concrete information put in to help CASTOR guys to debug.

Sites / Services round table:

- FNAL - switch crash about 05:00 local time. Network engineer rebooted. Back online within 2 hours. Same switch as crashed last week. Will replace some modules - a disruptive downtime (partial outage) later today.
- PIC - ntr
- KIT - yesterday asked when we will update FTS. Want to do it asap. RPMs still on ETICS - no official release nor release notes. Wait for them and then test in pre-prod then prod. Yesterday announced downtime for mainly storage - router to be updated and rebooted. No reply from experiments or TAB yet. WNs have been set offline. If VO does not want to lose any jobs should stop to submit jobs to GridKA - just storage will be affected.
- OSG - Retested alarm system to BNL and FNAL from GGUS yesterday. end to end ticket routing and SMS all successful. Request from an admin who got ticket forwarded via GGUS. Title was "SMS error" with no other details. Request resource name and brief description. Maybe an educational issue? Really helps to have this info in ticket description. Maria - do you have ticket # - Rob, can hunt it down. Maria - we have Twiki page with rules for submitting tickets. (MariaDZ) To answer RobQ Instructions for ALARMS in [https://twiki.cern.ch/twiki/bin/view/EGEE/SA1\\_USAG#Periodic\\_ALARM\\_ticket\\_testing\\_ru](https://twiki.cern.ch/twiki/bin/view/EGEE/SA1_USAG#Periodic_ALARM_ticket_testing_ru) (point 4.1. in case of full alarms where the whole chain should act). Please give us the ggus ticket number you mentioned.
- CNAF - ntr
- BNL - ntr
- ASGC - ntr
- NL-T1 - upgraded SRM this morning. More memory to fix problem ATLAS had. Now upgrading DPM to gLite 3.2. Also firmware upgrade for disk storage.

- RAL - problem this morning with CASTOR for CMS and ALICE. A DB node ran out of memory. Noticed and fixed - noone raised a ticket! At risk tomorrow for memory upgrade to CASTOR Oracle RAC
- GridPP - ntr
- IN2P3 - unsched outage yesterday afternoon - problem not really understood. Nearly all WNs got stuck in kernel module. No connections to WNs - had to reboot all WNs and lost about 9K jobs. Still are investigating. Up to date with kernels - suspect due to some of our specific s/w like AFS or GPFS or ... No idea for the moment of what caused this problem. Will file an SIR when we know abit more.
- NDGF - ntr
- CERN - reminder: CASTOR public in downtime tomorrow morning. Experiment services should be available.
- CERN DB - one node of pabr database has been since Sunday morning. Down due to clusterware problems - should be back in an hour or two.
  
- FTS - update on procedures. 2.2.3 still in staged rollout. Only possibility is to get to gLite preview page. Looks very much like standard release page. If you have 2.2.0 or 2.2.1. If you have 2.1 then follow procedure to 2.2.x. 2.0 (ASGC) have asked for advice for this site. Will require a complete new installation - different SL version. Documented procedure from 2.1 to 2.2.3 being tested. Should be available tomorrow or day after. Date for release to production: should be by first part of this week. Simone - please translate names of pages into Twiki URLs. Andrea - preparing email.

AOB:

## Wednesday

Attendance: local(Jamie, Miguel, Rosa Maria Garcia Rioja (Rosa), Edoardo, Harry, Oliver, Jean-Philippe, Nilo, Eva, Timjur, Simone, Roberto, Patricia, Jan);remote(Nicolo, Gonzalo Merino (PIC Tier1), Jon Bakken(FNAL), Angela Poschlad (KIT), anon, Michael Ernst(BNL), Gang Qin (ASGC), Rolf, Maria Dimou, Tiju Idiculla (RAL)), Onno Zweers (NL-T1), Vera Hansper (NDGF))

Experiments round table:

- ATLAS reports -
  1. SARA has shown instabilities (SRM) also this morning. Seems OK now. Onno - had some SRM problems as observed. Not stable at the moment. Think problem came from removal of other bottlenecks in whole architecture and now SRM is the bottleneck. Cannot handle it very well. Doubled internal memory - seems to help somewhat. Not sure if stable enough. Still investigating. Thanks for letting us know when things go wrong! Simone - did you observe new or increased activity. Onno - yes quite some activity in past two days. This probably exposes the instability.
  2. TRIUMF reported the completion of the migration to Chimera. Back in business.
  3. BNL has shown 1 hours of SRM problems during the night. Michael - observation was that these failures occured when there was a massive stage out from reprocessing jobs. >30K SRM requests. Queue depth not big enough - may want to increase it. Simone - bring up to ATLAS; a bit of a change in reprocessing schema. Simone - BNL and SARA hold 100% of data for reprocessing on disk.
- CMS reports - Highlights:
  - ◆ T1s
    1. MC reprocessing done at PIC, IN2P3 (except for a couple of jobs to be resubmitted.)
    2. Skimming tests at RAL, FNAL, IN2P3
    3. Full rate of backfill tests of CREAM CEs at European T1s.
  - ◆ ASGC
    1. Restarted backfill jobs at ASGC after setup of tape family.

- ◆ IN2P3
  1. Unscheduled storage downtime affected transfers, SRM now back.
- ◆ FNAL
  1. Site unreachable since 12:00
- ◆ CNAF
  1. Investigation timeouts in import of MC from Vienna.
- ◆ T2 highlights
  - ◇ MC ongoing in RAL, CNAF, PIC, FNAL T2 regions 1. Configured submission to NorduGrid, testing.
  - ◇ Ongoing SAM test failures at T2\_IN\_TIFR - DPM read errors. [ Now seem to have been solved. ]
- ALICE reports [↗](#) - GENERAL INFORMATION: following the report provided yesterday, the pass 1 and 2 reconstructions are still ongoing at T0 and some T1 (CNAF, FZK, RAL) with no remarkable issues to report. In addition there are 2 analysis trains also ongoing and no new MC cycle requests in the queues.
  - ◆ T0: Last week an overload of wms214.cern.ch was announced by Maarten. The problem disappeared in few hours with no external interventions. Waiting for the new massive MC production to follow the status of this WMS. At this moment, the number of active jobs is so low that the problem cannot be reproduced.
  - ◆ T1s:
    - ◇ CNAF: the new CREAM-CE entered production yesterday after the ops. meeting with no remarkable issues to report
    - ◇ CCIN2P3: Last week the site admin observed that ~1000 ALICE jobs were killed due to a too large memory consumption. The experiment agreed with the site that the site admins would allow the access of ALICE to a very long local queue. The queue seems to be ready but not yet published in the info system. The reason for this request is to enable that those jobs requiring a large memory (over 4GB) will be able to finish successfully.
- LHCb reports - DIRAC is switched off. No activity going on at all.

Sites / Services round table:

- PIC - ntr
- FNAL - right now we have significant building power issues. People working on it. Told to turn all CMS equipment off and keep it off. Until then we'll be down. All the info we currently have.
- KIT - tested FTS updates successfully in preprod. Will update production tomorrow during downtime
- CNAF - ntr
- BNL - ntr over what was said above. JPB - waiting for LFC logs!
- ASGC - ntr
- IN2P3 - SRM: had a h/w problem with machine running SRM. CMS noticed -currently in "at risk". Restarted machine and problem perhaps still there. Waiting for technical maintenance and are setting up a spare machine in case we have to take m/c out. Had 2 power cuts for remotely hosted WNs. About 170 WNs hosted 300km from CCIN2P3. Suffered from cut during night and a very short one this morning. Impacted ATLAS - about 170 jobs which crashed. Memory problem of ALICE jobs - working on this. Need to understand how we can still have capacity on WNs for other VOs if ALICE really needs >4GB. To solve this monitoring jobs and try to understand how much memory they really need. Not 4GB but a bit less... Adjusting local parameters for batch system to find a solution which allows us to optimize system. Harry - ALICE reported they have got down from 4GB to 3.1GB and soon back to 2GB. (LHCC refs meeting).
- RAL - ntr
- NL-T1 - can elaborate on SARA SRM: in the last half year we have had made some architectural changes. Storage - compute interconnect now 40Gb. Increased # storage nodes and # WNs. SRM is

on list to be upgraded. New h/w there but not ready yet. SRM is the weak link at the moment. Had some load in last few days from FTS and from compute cluster. Not only ATLAS but also a few other VOs. Combination of load from several sources. Noticed on Monday that SRM was swapping hence increase of memory from 4GB to 8GB. seems to help somewhat but not sufficient as still problems. Restarted dCache this morning. One of services did not start up ok - restarted again and then ok. New h/w that is available as 24GB of memory so swapping should not be a problem. Seems that swapping was not only issue...

- NDGF - ntr
- OSG - ntr
  
- CERN DB - ATLAS replication to BNL currently stopped. Importing missing data from CERN to them. Cause of this missing data has been identified by Oracle. Part of Streams dictionary missing at dest. Working to fix. ATLAS replication to KIT disabled this morning. Streams admin p/w changed and ATLAS DBAs use a monitoring tool that uses this account (Should not!) and blocked this account. (Another acct already requested).
  
- FTS upgrades. 2.2.3 released to rollout and is recommended version. Upgrades from 2.2.0 ok. Sites on 2.1. Checked scripts. Should be ok. PIC have done this(?) If any site wants to do upgrade propose for first site to do this together with FTS SUPPORT. Simone - if sites upgrade please let me know so we can check checksum verification. Gonzalo - upgrading a test instance. Testing procedure and functionality. Finding some issues and in contact with support through GGUS. Oliver - issues with checksums being investigated.

AOB: (MariaDZ) The ticket that RobQ mentioned yesterday was a TEAM and not an ALARM ticket. GGUS documentation is in <https://gus.fzk.de/pages/docu.php#team> As this was an ATLAS ticket there is more advice to submitters in [https://twiki.cern.ch/twiki/bin/view/Atlas/ADCPPoint1Shift#Submitting\\_a\\_GGUS\\_Ticket](https://twiki.cern.ch/twiki/bin/view/Atlas/ADCPPoint1Shift#Submitting_a_GGUS_Ticket) As the site=resource was preselected at submission time in GGUS (see [https://gus.fzk.de/ws/ticket\\_info.php?ticket=55426](https://gus.fzk.de/ws/ticket_info.php?ticket=55426)) please say what more should have been included or if anything wasn't mapped into the OSG ticketing system. If development is needed please follow up via savannah. Rob - will check and respond.

## Thursday

Attendance: local(Jamie, Dirk, Timur, Julia, Jan, TIm, Harry, Andrea, Jean-Philippe, Lola, Eva, Nilo, Simone, Roberto, MariaDZ);remote(Jon/FNAL, Gonzalo/PIC, Rolf/IN2P3, Angela/KIT, cnaf, Gang/ASGC, Michael/BNL, Ronald/NL-T1, Alexei, Vera/NDGF, Nicolo).

Experiments round table:

- ATLAS reports -
  1. Data re-processing (Nov09-Dec09 data). The reconstruction step is done (RAW->ESD production). Merging step is in progress. Merged HIST datasets are sent to CERN for validation.
  2. LYON : unscheduled downtime for 1.5h
  3. ASGC : ANALYSIS QUEUES status is set to 'broker off' (Don't send user analysis jobs to Taiwan for the moment but continue to send test jobs).
  
- CMS reports - Highlights:
  - ◆ CERN
    1. Recovering backlog of transfers P5 --> T0 (being recovered).
    2. node from cmsexpress queue in maintenance (high load on node)
  - ◆ T1s
    1. Full rate of backfill tests of CREAM CEs at European T1s.
  - ◆ IN2P3

1. New unscheduled SRM downtime (affected transfers and other activities - also reported by site contact).
  - ◆ FNAL
    1. Recovering from power outage - all services available, dCache disk pools running at reduced capacity for safety reason.
  - ◆ CNAF
    1. Timeouts in import of MC from Vienna - Vienna admin reset NIC on gridftp server, files transferred. (i.e. problem solved).
  - ◆ Services
    1. Gridmap unavailable. (currently not displaying any data). Julia - application supported by monitoring team in IT-GT. Made an upgrade recently which is probable cause. Tim - reason for backlog? Nicolo - will check: transfer system on P5 side. During weekend or Monday.
  - ◆ T2 highlights: MC ongoing in IN2P3, CNAF, PIC T2 regions
    1. T2\_US\_MIT, T2\_IT\_Rome, T2\_ES\_CIEMAT scheduled downtime.
- ALICE reports [↗](#) - GENERAL INFORMATION: at this moment there are no new MC cycles in production nor analysis trains. The jobs currently running belong to the Pass1 reconstruction announced yesterday
  - ◆ T0 site - Pass 1 reconstruction currently running using CREAM-CE resources with no incidents to report
  - ◆ T1 sites - Good behaviour of the T1 sites that has executed the Pass 2 reconstruction finished yesterday
- LHCb reports - DIRAC has been restarted yesterday according the schedule. Few minor problems reported by users due to a changed interface and GANGA stick to old version of DIRAC. Promptly done the alignment GANGA/DIRAC. Launched a small production yesterday evening to test the new brand of setup and no major problems have been reported with just a few percent of jobs failing. SAM activity wasn't impacted too much while other information on SLS and Dashboard had to be adapted to the new DIRAC. There are currently in the system few thousand jobs running for some small MC productions (6M-12M events to be produced each).
  - ◆ T2 sites issues:
    - ◇ INFN-LHCb-T2: issue with permissions in the shared area
    - ◇ CESSGA: issues uploading logs/sandbox to LogSE
    - ◇ Padova, Sofia and Bologna: issue with g++ compiler missing. Worth to mention explicitly in the VO-ID Card

#### Sites / Services round table:

- FNAL - CMS services have been restored by 7 PM Chicago time. For safety reasons, we were allowed to bring up 75% of our previous power load in the affected building. Our plan did this by only reducing the online dCache pools - that means all interactive access, BlueArc data, and condor job capabilities are at 100%. The downside is that data service is slower for files not on the powered disks since they have to be restored from tape. The power situation continues to be evaluated by experts and hopefully by Thursday PM we will have a plan for the 25% that are currently powered off. There isn't any planned outage for Thursday. Hope to be back 100% by end of the day.
- PIC - ntr
- KIT - in downtime. Looks good so far!
- CNAF - ntr
- ASGC - CMS test jobs running fine at ASGC. When they finish will online 20 tape drives and some disk servers.
- BNL - ntr
- NL-T1 - SARA had problem with one of dCache head nodes /var filled up. Fixed and running ok. Tuesday: LFC for LHCb and ATLAS will be moved to new h/w. Short. FTS & FTA will also be

moved at same time.

- NDGF - ntr
- RAL - We have two scheduled outages in the GOC DB.
  1. On Tuesday 23rd February, there is an 'AT RISK' on all castor instances due to an NFS reconfiguration.
  2. From tomorrow (19th Feb) to Tuesday (23th Feb) there is an 'OUTAGE' on lcge07. This is to allow for a drain and a disk replacement.

#### APEL status update

the APEL central database is now back online, and data are updated on the accounting portal. The latest data appearing there are from Monday 15th Feb. Recently published data will pop up regularly now. Sync pages and SAM tests are updated as well. After our integrity checks, we can guarantee that 99.8% of the data have been restored, and we have no reason to think that the remaining 0.2% have been lost. We however advise all sites to check their data on the accounting portal and report any inconsistencies through a GGUS ticket. A daily update of APEL status can be obtained from the following page:[http://goc.grid.sinica.edu.tw/gocwiki/ApelIssues-Jan\\_Feb\\_2010](http://goc.grid.sinica.edu.tw/gocwiki/ApelIssues-Jan_Feb_2010) A broadcast has been sent to the EGEE community. Publication is working fine, but APEL is still to be considered AT RISK until Monday the 22nd. More information can be obtained by mailing [apel-support@jiscmailNOSPAMPLEASE.ac.uk](mailto:apel-support@jiscmailNOSPAMPLEASE.ac.uk)

- IN2P3 - update on SRM issue. Yesterday we had h/w problem with SRM server. The consequence was firmware update which explains unscheduled downtime. Should be ok now - investigating what we can do to improve failover. Downtime on March 2nd for maintenance on HPSS and BQS. Outage of ~6 h 08:00 - 14:00. Batch system concerned hence long running jobs will be blocked the day before. Normal sequence of downtime announcements will occur.
- CERN DB - problem with ATLAS offlien DB last night. One disk group running out of space - no connections available to DB. Have improved monitoring to be more proactive for such problems. BNL - ATLAS replication completely resynchronized.
- CERN - repeat of ATLAS xrootd requests to CASTOR block - affected many users. Have fix - change assessment submitted to ATLAS (Simone) for approval. Could be put in tomorrow even though its Friday. Simone - sent mail to ATLAS saying its scheduled for tomorrow at 10:00 unless serious complaints by 17:00 today. Change assessment: <https://twiki.cern.ch/twiki/bin/view/CASTORService/ChangescastoratlasXrootd19Feb2010> (aka "risk assessment").
- CERN SLS for CASTOR Monitoring - minor issue. Upgrade of problem - some instances marked as grey. Fixed since.

AOB:

- One but last USAG meeting prior to end EGEE. Agenda: <http://indico.cern.ch/conferenceDisplay.py?confId=85636>. Click on Doodle (in this agenda) if you would like to propose a different date. **NB!** Meeting at 4pm (KV time zones and summer time!) to suit participants from the Americas.

## Friday

Attendance: local(Jamie, Andrea, Jean-Philippe, Lola, Timur, Gavin, Ueda, Patricia, Simone, Malik);remote(Jon/FNAL, Xavier/KIT, Michel/BNL, Rolf/IN2P3, Gang/ASGC, Tiju/RAL, Rob/OSG, Onno/NL-T1, Alexei/ATLAS, CNAF, anon (Brian?), Roger Oscarsson (NDGF)).

Experiments round table:

- ATLAS reports -
  1. Ultra-urgent simulation tasks issue (more info after daily ADC phonecon)
  2. Reprocessing Coordinator proposed updated ESD replication policy

3. Oracle database outage (see Feb 18). Luca Canali provided post-mortem (thanks for this - very useful!)

- CMS reports -
  - ◆ T1s
    1. Pre-production for next round of ReReco and Skimming at PIC, CNAF, FNAL.
      - 52 jobs stuck in 'Scheduled' state at CNAF - fixed by site admins
    2. Prompt skimming tests at FNAL, RAL, IN2P3 & KIT.
    3. Low rate of backfill test jobs of CREAM CEs at European T1s.
  - ◆ IN2P3
    1. Reduced quality on T0-->IN2P3, GGUS #55707 (Rolf - reduced transfer quality due to combined load from CMS & ATLAS on gridftp servers. Increased # of servers.
  - ◆ RAL
    1. 20 files on T0-->T1\_UK\_RAL not transferring since last Friday - probably issue with PhEDEx Download agent at RAL.
  - ◆ ASGC
    1. Good results from first round of backfill jobs, waiting for OK from site admins before submitting second round.
    2. One corrupted file on tape at ASGC - retransferred from source T2.
  - ◆ T2 highlights
    - ◇ MC ongoing in KIT, IN2P3, CNAF T2 regions
    - ◇ SAM test errors at T2\_BE\_UCL on CE: Unspecified gridmanager error - Job got an error while in the CondorG queue.)
    - ◇ SAM test errors at T2\_BR\_SPRACE (Maradona error on CE + host cert expired on SRM)
- ALICE reports [↗](#) - GENERAL INFORMATION: During yesterday's ALICE TF meeting, the new MC cycles expected before the next data-taking in March were discussed. A new MC cycle might run before data-taking with a low number of jobs (final decision still pending). In the meantime, the pass1 reconstruction at the T0 is still ongoing and in addition, there is one analysis train in execution. Last night at 04:00 CET the air conditioning in the AliEn central services server room failed. Subsequently, the central services are unavailable. The experts are fixing the problem now.
  - ◆ T0 site - Few jobs are currently running at the T0 with no incidents to report
  - ◆ T1 sites - Production stopped at all T1 sites
- LHCb reports - There are no large activities on going in the system. Mainly user problems due to a non completely transparent migration of DIRAC backend (proxies not uploaded and then jobs failing with proxy expiration).
  - ◆ SLS reported for all T1's and T0 read-only a shortage this morning at about ~2:00 UTC. Perhaps connected with some known streams replication issue? (see detailed report)
  - ◆ T1 sites issues: GridKA: SQLite problem
  - ◆ T2 sites issues: Shared area problem at Barcelona

Sites / Services round table:

- FNAL - ntr
- KIT - Yesterday's intervention(s) successfully finished
- BNL - ntr
- IN2P3 - nothing additional
- ASGC - ntr - onlining tapes and disk servers
- RAL - ntr
- OSG - outage at ATLAS Great Lakes T2 - storage experts on it. GGUS ticket submitted Wednesday turned into BNL ticket 15463 BNL->Dallas site.
- NL-T1 - SRM problems at SARA: earlier this week we reported that we fixed swapping of SRM node

and yesterday reported one of dCache nodes /var was full. Today we fixed 2 issues: 1 pool node on which dCache had crashed. Restarted service and running fine. Issue with reading files from tape - this has to do with migration from pnfs to Chimera. Did this in January and files that have been written since to tape could not be restored from tape due to bug in script. Changed script to read and write to tape after Chimera migration. Fixed bug this morning and reading from tape should now work. Issues still not solved - graph showing many failed jobs. Don't have more info at this time. Would welcome further info from ATLAS preferably in GGUS ticket 55643

- CNAF - ntr
- NDGF - ntr
  
- **CERN LHC VOMS Service** : For the LHC voms service a routine update of the host certificate for lcg-voms.cern.ch will happen next week. By 09:00 UTC on 25th of February gLite 3.1 services requiring lcg-vomscerts should be updated. They should be updated to the now released lcg-vomscerts-5.8.0-1.noarch.rpm. A reminder will be posted in this meeting the day before as well as general broadcast early next week.
  
- CERN WMS now reconfigured to support all requested experiment roles.
  
- CERN - glitches in xroot in CASTOR ATLAS upgraded xroot ssl module with fix from developer. Seems to fix problem.

AOB:

-- JamieShiers - 11-Feb-2010

---

This topic: LCG > WLCGDailyMeetingsWeek100215

Topic revision: r17 - 2010-02-19 - JamieShiers



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