

# Table of Contents

<b>Week of 110131</b> .....	<b>1</b>
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday:.....	1
Tuesday:.....	3
Wednesday.....	5
Thursday.....	8
Friday.....	10

# Week of 110131

## 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](#)

## WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments

VO Summaries of Site Usability				SIRs, Open Issues & Broadcasts			Change assessments
<a href="#">ALICE</a>	<a href="#">ATLAS</a>	<a href="#">CMS</a>	<a href="#">LHCb</a>	WLCG Service Incident Reports	WLCG Service Open Issues	Broadcast archive	CASTOR Change Assessments

## General Information

General Information				GGUS Information	LHC Machine Information
CERN IT status board	M/W PPSCoordinationWorkLog	WLCG Baseline Versions	WLCG Blogs	GgusInformation	Sharepoint site - Cooldown Status - News

## Monday:

Attendance: local(Alessandro, Oliver, Jarka, Maria, Jamie, Maarten, Stefan, Andrea, Gavin, Ewan, Jan, Mike, Roberto, Dirk, Eva);remote(Ron, Jon, CNAF, Gonzalo, Michael, Dimitri, Rolf, Gareth, Suijan, Rob, Daniele, Christian Sottrup).

Experiments round table:

- ATLAS reports -
  - ◆ ATLAS has new data project tags, the data11\_7TeV and data11\_900GeV
  - ◆ PIC provided list of lost files; GGUS:66409
  - ◆ SARA-MATRIX: storage degraded saturday morning. TEAM ticket escalated to ALARM (see below). SRM services restarted, issue fixed.
  - ◆ TEAM ticket GGUS:66766 escalated to ALARM, however, ALARM was rejected by site because the TEAM ticket originator does not belong to alarmers. AMOD then copied the 66766 into new TEAM ticket GGUS:66768, this has then been escalated to ALARM with success (TEAM ticket originator is one of the alarmers). It seems that the escalation works only if the submitter is both TEAMer and ALARMer at the same time.
- CMS reports - NTR
- ALICE reports -

- ◆ T0 site
  - ◇ Saturday there was a problem with voalice13 (WMS submission for ALICE and PackMan server), the internal disk broke. voalice14 has entered in production to substitute it and it is performing well. We asked for a replacement of voalice13 disk. This took a few hours. AliEn version at CERN had some features that had to be configured for new VO box.
  - ◇ CAF: problem accessing OCDB on AAF since the end of last week was solved by downgrading AliEn version to 2.18.
- ◆ T1 sites
  - ◇ KIT: cream-1, cream-3 and ce-3 have been put in an unscheduled downtime this morning. It is scheduled until 16:30. There were no jobs running since Saturday, the issue seems to be caused by license server problems
  - ◇ RAL: ticket GGUS:66711 [↗](#), problem identified but ticket still open. Any news about this? [ Gareth - ticket will be updated soon. Puzzling: things that were found have been there for quite some time. Things seem to have been at fault longer than problem - will continue looking ]
- ◆ T2 sites
  - ◇ Usual operations
- LHCb reports - MC productions. Re-stripping. General overload on almost all SRM services caused by demanding stripping jobs. Reduced the concurrent number of jobs running in parallel on each site.
  - ◆ T0
    - ◇ CASTOR vs xroot on Friday (GGUS:66731 [↗](#)). The xrootd server crashed. Now it is OK ticket left open [ Jan - s/w change to be deployed by LHCb to avoid using default service class. Is this being rolled out? R - patch in Gaudi but don't know status of deployment. ]
    - ◇ VOBOX: following the scheduled intervention this morning some update screwed up the ssh port and not login was possible. Fixed.
  - ◆ T1
    - ◇ IN2P3 Many job failures on Saturday opening files via dcap or resolving turl (GGUS:66745 [↗](#)). The ticket hasn't been updated since Friday. [ Rolf - will follow up with ticket ]
    - ◇ RAL: Scheduled downtime for Oracle upgrade. Site drained since Sunday.
    - ◇ PIC: Pilot jobs aborting on the CREAMCE (GGUS:66793 [↗](#))
    - ◇ SARA: issue on Saturday with SRM that was down and restarted few hours later. Discussion in a thread no GGUS opened. Backlog of re-stripping jobs formed over there.
  - ◆ T2 site issues: :
    - ◇ A lot of problems at non-T1 sites with ~10 tickets opened since Friday

Sites / Services round table:

- NL-T1 - as mentioned by LHCb and ATLAS Saturday morning SRM crashed. Restarted Saturday afternoon with increased limits - ran into some memory limit. This seems to have solved issue.
- FNAL - ntr
- CNAF - ntr
- PIC - brief update on issue with ATLAS data. Situation as of today: on Friday 60% of affected files put online - dCache saw pools and files visible through dCache. A big number of errors seen in transfers no longer there. Files still on partially damaged pools. Migrating to different h/w. Other 40% copying outside of dCache into different h/w. Will be able to recover most data which will take some days. # of really damaged files will be about 1000. Until we copy all files we will not have list of files affected. The list of files in these pools made available to ATLAS on Friday. Was an intervention in a scheduled downtime to replace two of four controllers of the disk system. In the process of this substitution one of the controllers had a newer firmware version. Never put online. Looks as if for a

few ms was online and could create this corruption. e.g. 2 servers of SAN seeing same disks. No evidence of this in logs. Investigation still ongoing. A SIR will be produced when the incident is closed.

- BNL - ntr
- KIT - problem with one of our sub clusters. Seems to be problem with license server. Extended downtime until tomorrow afternoon as we need support from developers and this will take time. Move some WNs to another cluster - supports 50% of WNs up from 30%. Will also include one ore CREAM CE to support more jobs submitted to the working cluster.
- IN2P3 - nta
- RAL - outage that we are now in for Oracle updates to CASTOR DBs is going ok. Scheduled to get CASTOR back during the afternoon. A longer outage for CMS to update disk servers to 64bit O/S and then enable check-summing. Need to schedule updates of non-CASTOR DBs for ATLAS and LHCb - possibly Tue/Wed next week.
- ASGC - ntr
- NDGF - ntr
- OSG - ntr
  
- CERN storage - one team ticket from ATLAS about SRM transfers timing out - faulty disk server in a bad shape only affecting LSF. Once restarted ok. Jarka - downtime of CASTOR - how long does it take? ~3 hours.

AOB:

## Tuesday:

Attendance: local(Luca, Roberto, Gavin, Jamie, Maria, Mattia, Mike, Lola, Nico, Stefan, Ewan, MariaDZ);remote(Michael, Francesco, Kyle, Ronald, Gareth, Jeremy, Jon, Paolo, Elena, Xavier, Suijan, Elena, Christian, Gonzalo).

Experiments round table:

- ATLAS reports -
  - ◆ PIC GGUS:66409: 60% of affected files were put online; pic and T2s in ES cloud are running production and analysis jobs since yesterday afternoon.
  - ◆ FZK has reduced performance (FZK is in downtime, license server problems of the batch system). [ Xavier - PBS of one of subclusters broken and had to be re-installed from scratch. Will be finished within downtime - all should be working by 19:00 ]
- CMS reports -
  - ◆ CERN
    - ◇ No T0 activity today
      - First 2011 MidWeekGlobalRun Tue 2nd-3rd
      - Heavy Ion zero-suppression and reconstruction possibly starting this week
    - ◇ Several reports from operators of ssh errors/freezes on lxplus, vocms\*, svn, trac possibly related to Kerberos intervention, looks OK since 12:30 CET so no ticket submitted
      - <http://it-support.servicestatus.web.cern.ch/it-support-servicestatus/ServiceCha>
      - <http://it-support.servicestatus.web.cern.ch/it-support-servicestatus/110201AFS>  
[ Gavin - around this time saw AFS problems triggered by kerberos intervention. Load related? Remains to be understood. ]
  - ◆ T1s
    - ◇ MC reprocessing with 3\_9\_7 is in the tails.

- RAL - scheduled downtime for CASTOR upgrade (checksumming activation on CMS disk servers)
  - KIT - unscheduled downtime of CREAM
- ◆ T2s
  - ◇ Spring11 8TeV production was running at many sites (also at Tier-1s)
    - Note: CMS will decide this afternoon what to do with the ~400M 8TeV events already produced.
  - ◇ T2\_DE\_DESY UNSCHEDULED downtime (issues with the scheduled firewall reconfiguration)
- ◆ Other
  - ◇ GGUS briefly unavailable around 12:00 CET for network glitch at KIT, now back.
- ALICE reports [↗](#) - mid Feb start reprocessing of HI run
  - ◆ T0 site
    - ◇ Nothing to report
  - ◆ T1 sites
    - ◇ RAL: ticket GGUS:66711 [↗](#), problem solved
  - ◆ T2 sites
    - ◇ Usual operations
- LHCb reports - MC productions running smoothly almost everywhere. Re-stripping: tail of jobs remaining at SARA
  - ◆ T0
    - ◇ Sent a request to relevant CASTOR people to help our Data Manager to understand/reconstruct why some files present in high level catalogs are not available physically on the Storages (not only at CERN)
  - ◆ T1
    - ◇ IN2p3 Many job failures on Saturday opening files been understood: files were garbage collected and tnot longer available on the disk cache.
    - ◇ RAL: Data integrity check preventing some stub files to be migrated to tape (spotted order of tens of such corrupted files per day). Proposed to Shaun to use the same script currently used by CASTOR at CERN plus mail notification to data manager to cross check all these files. [ Gareth - plan is tomorrow to roll-out an updated gridftp component of CASTOR which will user 2.1.9-10 version that fixes this problem, declare an AT RISK tomorrow to roll-out this change on LHCb disk servers. (Checksums created will agree with file as written which will allow migrations to tape to proceed normally)]
    - ◇ RAL: All MC jobs failing at RAL with pilot not running any longer. GGUS:66849 [↗](#) [ Gareth - a couple of problems on batch system - drained Sunday ready for CASTOR outage yesterday. Due to oversight some jobs started middle of yesterday which were then paused. Also had config problem on batch system - the prologue would run then job fail, be re-queued etc. Largely sorted out this morning - batch now running ok. Problem cause of MC problems. ]
    - ◇ SARA: draining the remaining re-stripping jobs accumulated over there. The process goes very slowly due to the limited Fair Share and MC jobs competing with.
    - ◇ GridKA: almost all farm is available now (after the license problem in some CEs) and we observed indeed a net increase on the number of jobs running there.
  - ◆ T2 site issues: :
    - ◇ NTR

Sites / Services round table:

- IN2P3:
  - ◆ LHCb mentioned yesterday a GGUS ticket

Tuesday:

([https://gus.fzk.de/ws/ticket\\_info.php?ticket=66745](https://gus.fzk.de/ws/ticket_info.php?ticket=66745)) opened on Friday and apparently without response on Monday. After investigation it appeared that an answer had been given in the site's local help system but because of a temporary problem in the interface to GGUS this action was not reflected centrally. In the meantime LHCb suggested to close the incident.

- ◆ Another point is a pre-announcement of an all day outage next week on Tuesday, February 8th. This involves the following tier1 and tier2 services: batch and interactive, tape mass storage (HPSS+robotics), dCache local network, Oracle and all Oracle based services like LFC for example, EGI operations portal.
- ◆ Because of the outage of the Operations portal, **all** notifications for downtimes of EGI sites will be delayed during the time of inavailability of the portal. However, as the GOCDB is hosted elsewhere (at RAL) and contains a complete list of all declared outages, one can still get a list by using the GOCDB's interface.
- ◆ For clarification, as already announced there will probably be another outage of the tape mass storage system later this month. An update on this will be given one week before, as usual.

- BNL - ntr
- CNAF - ntr
- RAL - problem with a tape - will contact LHCb via GGUS ticket. Update to Oracle DB went ok. Update to CMS disk servers to 64bit OS means we can enable checksums for CMS and will do this for one storage class tomorrow. Propose Wed next week to update other Oracle DBs to 10.2.0.5.
- FNAL - ntr
- KIT - ntr
- NL-T1: downtime announcement for NIKHET on 15 Feb. Entire site down to replace core router. Old one has to be disconnected and removed before new one can be installed. At risk downtime for 2 days adjacent to this to hunt down and fix any problems.
- ASGC - ntr
- NDGF - reminder of AT RISK tomorrow afternoon for ATLAS data.
- PIC - ATLAS data corruption: copying out of damaged f/s on-going. No further news over yesterday - last update to ticket yesterday morning. ATLAS was going to produce a list of files from their point of view. Still being produced? [ Elena - will check ]
  
- OSG - ntr
- GridPP - ntr
  
- CERN - ntr

AOB: (MariaDZ) About the ATLAS remark on TEAM tickets being upgradable to ALARM ones only by a TEAMer who is Also ALARMer, please remember that this was part of the original description of what GGUS would implement. The sites insisted from the beginning to keep the list of ALARMers short. This is why not all TEAMers should be able to upgrade tickets into ALARMs. Nevertheless, TEAM tickets belonging to All TEAMers, those with appropriate rights should be doing the upgrade. All this was presented at the daily meeting <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGDailyMeetingsWeek101115#Wednesday> and the 20101123 MB. The description was recorded in <https://savannah.cern.ch/support/?118153#comment0>

## Wednesday

Attendance: local(Andrea, Carles, Eva, Ewan, Maarten, Maria D, Mattia, Nicolo, Roberto, Stefan);remote(Christian, Daniele, Dimitri, Elena, Francesco, Gonzalo, Jeremy, Jon, Onno, Rob, Rolf, Suijian, Tiju).

Experiments round table:

- ATLAS reports -

- ◆ T0/CERN, Central services
  - ◇ ATLAS elog was unavailable for 20 min around midnight.
- ◆ T1s
  - ◇ PIC GGUS:66409: production and analysis jobs are running at pic and T2s in ES cloud. ATLAS provided a list of high priority datasets.
  - ◇ FZK GGUS:66783: FZK was set fully online after DT last evening. There is a network problem to communicate with servers CERN (only 20% of jobs are affected).
- CMS reports -
  - ◆ CERN
    - ◇ Restarting T0 activity
      - MWGR starting at P5, subdetectors joining gradually. T0 Ready to receive data, will not be exported to T1s.
      - Heavy Ion data pre-staged from tape in preparation for zero-suppression.
    - ◇ Another service used by CMS affected by Kerberos upgrade yesterday: CVS, fixed around 17:00 CET.
      - <https://remedy01.cern.ch/cgi-bin/consult.cgi?caseid=CT0000000742358&email=magi>
      - In general, since so many services depend on Kerberos, CMS would appreciate the declaration of an "AT RISK" downtime in case of such interventions.
        - Ewan: will have the matter followed up
    - ◇ CMS Elog unavailable between 23:30-24:00 CET, came back by itself.
      - GGUS:66887 was opened during the night, can be closed as soon as issue is understood.
  - ◆ T1s
    - ◇ MC reprocessing with 3\_9\_7 is in the tails.
    - ◇ ASGC - 25 TB of data produced at ASGC in Dec22nd reprocessing campaign was not migrating to tape for custodial storage
      - GGUS:66898 - thanks to ASGC for intervention on Chinese New Year's day.
  - ◆ T2s
    - ◇ Spring11 8TeV production stopping; completing currently running jobs.
    - ◇ T2\_ES\_IFCA - scheduled downtime
    - ◇ T2\_PK\_NCP - unscheduled network outage
    - ◇ T2\_US\_Vanderbilt - Vanderbilt upgraded from T3 to T2 (dedicated to Heavy Ion processing), started CMS Site Commissioning procedure
      - Savannah:118948
- ALICE reports -
  - ◆ T0 site
    - ◇ CAF nodes: reinstalled the nodes to modify the partitioning.
    - ◇ There were no jobs running on CREAM CE due to a problem of permissions in new ALiEn version. Instead of passing the AliEn user from a configuration file to the working node, it was trying to pass the VOBOX local user (alicesgm)
    - ◇ GGUS:66900: checking the information published for CERN CEs (lcg-infosites), we found some strange values: it reported that we had 0 CPUs at CERN. The information was very unstable as it goes from time to time from 0 CPUs to thousands. Update: expected behavior if the batch system becomes unresponsive, which has happened several times in last couple of weeks. The experts are looking at it.
  - ◆ T1 sites
    - ◇ Nothing to report
  - ◆ T2 sites
    - ◇ Usual operations

- LHCb reports -
  - ◆ Experiment activities:
    - ◇ MC productions running smoothly almost everywhere.
  - ◆ New GGUS (or RT) tickets:
    - ◇ T0: 0
    - ◇ T1: 2
    - ◇ T2: 0
  - ◆ Issues at the sites and services
    - ◇ T0
      - Sent a request to relevant CASTOR people to help our Data Manager to understand/reconstruct why some files present in high level catalogs are not available physically on the Storages (not only at CERN)
    - ◇ T1
      - RAL: Transparent intervention for LHCb CASTOR to upgrade the gridftp server using the right checksum
        - Tiju: intervention completed, new situation being monitored
      - NIKHEF: many user jobs failing. It seems to be a data access problem. (GGUS:66287 [↗](#))
      - GridKA: A lot of pilot jobs failing there through CREAMCE (GGUS:66899 [↗](#)).
        - Dimitri: the CREAM CE were published incorrectly; fixed, but no new SAM results yet
    - ◇ T2 site issues:
      - NTR

Sites / Services round table:

- ASGC
  - ◆ reduced coverage because of Chinese New Year's holidays Feb 2-6, but a shift rota has been put in place
    - ◇ Happy New Year!
- CNAF - ntr
- FNAL
  - ◆ CMS services OK, but FNAL closed due to bad weather, which may delay response to tickets
- GridPP
  - ◆ status of SLC 5.6 backward incompatibility reported elsewhere by LHCb?
    - ◇ Ewan: Steve Traylen thinks there is just a single library to be recompiled
    - ◇ Stefan: tcmalloc
- IN2P3 - ntr
- KIT
  - ◆ subcluster that was unavailable is almost completely back; it was unstable after reinstallation, sometimes did not accept new jobs for a while
- NDGF
  - ◆ scheduled downtime proceeding as planned
- NLT1
  - ◆ LHCb ticket has been updated, we can supply more information as needed
- OSG
  - ◆ looking into hosting top-level BDII service on behalf of BNL and FNAL
  - ◆ weather conditions at the OSG GOC may slow down response to tickets
- PIC
  - ◆ ATLAS data recovery proceeding well, a list of high-priority file name patterns has been received
    - ◇ Elena: are the patterns OK or should actual file names be provided?
    - ◇ Gonzalo: patterns probably enough



- RAL - nta
- dashboards - ntr
- databases - ntr
- grid services - ntr
- networks - ntr

AOB: (MariaDZ) Test of the GGUS-SNOW interface is now open as requested at the 2011/01/25 presentation (notes in <https://savannah.cern.ch/support/index.php?118062#comment52>). Open a GGUS ticket in the training system <https://iwrgrustrain.fzk.de/pages/ticket.php> If you can't please register via <https://iwrgrustrain.fzk.de/admin/register.php>. Open the SNOW dev instance <https://cerndev.service-now.com/> and find your ticket. If you aren't authorised to use SNOW dev, please ask Zhechka. Any comments, please write to [ggus-info@cern.ch](mailto:ggus-info@cern.ch) (only loss of functionality please at this stage, no new desired features).

## Thursday

Attendance: local(Eddie, Jacek, Jan, Maarten, Manuel, Nicolo, Roberto, Stefan);remote(Christian, Daniele, Elena, Elizabeth, Foued, Francesco, Gareth, Gonzalo, Jon, Paolo, Rolf, Ronald, Suijian).

Experiments round table:

- ATLAS reports -
  - ◆ T1s ongoing issues
    - ◇ PIC GGUS:66409: 30% of lost files need to be recovered.
    - ◇ FZK GGUS:66783: 20% of atlas jobs are failing because of network problem.
- CMS reports -
  - ◆ CERN
    - ◇ Restarted T0 activity
      - MWGR at P5, including combined CMS-TOTEM run.
      - T0 Receiving data, not exported to T1s.
      - T0 prompt processing not fully working (application crashes + T0 manager issues).
  - ◆ T1s
    - ◇ MC reprocessing with 3\_9\_7 is in the tails.
    - ◇ ASGC - GGUS:66898 closed, tape migration resumed.
    - ◇ RAL - scheduled AT RISK downtime of SRM for disk server reconfiguration.
  - ◆ T2s
    - ◇ 7TeV production ongoing.
    - ◇ T2\_DE\_DESY - at risk for network maint
    - ◇ T2\_PT\_NCG\_Lisbon - scheduled downtime
    - ◇ T2\_EE\_Estonia - missing os rpms, request to install HEP\_OSlibs\_SL5 meta-rpm SAV:119020
    - ◇ Some MC files lost before custodial transfer to T1s, invalidating SAV:118784 and SAV:119033
- ALICE reports -
  - ◆ T0 site
    - ◇ CAF nodes: reinstall the nodes to modify the partitioning. We ran into an Anaconda bug that prevented us from doing the desired partitioning in three of the machines. Under investigation.
  - ◆ T1 sites
    - ◇ Nothing to report

- ◆ T2 sites
  - ◇ Cybersar: GGUS:66910. There are no jobs running since last week due to expiration of the admin's proxy.
- LHCb reports -
  - ◆ Experiment activities:
    - ◇ MC productions running smoothly almost everywhere.
    - ◇ tcmalloc library issue: it has been discovered an incompatibility (bug) of some tcmalloc methods (getCpu()) to work with the new kernel of RHEL5.6 (the basis of SLC5.6). A patch in Gaudi has been created to use instead a minimal version of this library that does not use this method and then bypasses this incompatibility. For information the bug is tracked here: tcmalloc bug.
  - ◆ New GGUS (or RT) tickets:
    - ◇ T0: 0
    - ◇ T1: 2
    - ◇ T2: 0
  - ◆ Issues at the sites and services
    - ◇ T0
      - Sent a request to relevant CASTOR people to help our Data Manager to understand/reconstruct why some files present in high level catalogs are not available physically on the Storages (not only at CERN)
    - ◇ T1
      - RAL: we lost some data on their tape. This data is available at CERN and will be replicated (GGUS:66853)
      - SARA: many user jobs failing. It seems related to a file access problem. (GGUS:66287)
      - SARA: Unscheduled restart of 4 dcache pool nodes for LHCb.
      - GridKA: also (but in less extent than SARA) user jobs failure rate a bit larger than usual. This is due to the high number of concurrent jobs running and accessing the storage creating an heavy load.
        - Roberto: correlation with increased number of concurrent user jobs at KIT and NLT1, which have been configured to allow such jobs
    - ◇ T2 site issues:
      - NTR

Sites / Services round table:

- ASGC - ntr
- CNAF - ntr
- FNAL
  - ◆ what DB version is used for the FTS at CERN?
    - ◇ 10.2.0.5
  - ◆ there is some pressure for FNAL to upgrade to 11g by July because the corresponding support contract would be a lot cheaper, but has the FTS been tested against that version?
    - ◇ Jacek: it should work, but has not been tested
    - ◇ Maarten: will follow up, more news at the next T1 coordination meeting
- IN2P3
  - ◆ a broken network link caused ~10% of our WN located in Montpellier to be unreachable since ~noon; the corresponding jobs are probably lost; being investigated by network experts and provider
- KIT
  - ◆ problem with job submission to 1 subcluster due to license server issue has been fixed now
- NDGF - ntr
- NLT1 - ntr

- OSG
  - ◆ top-level BDII deployment plans for USATLAS (GGUS:66866 [↗](#)) and USCMS (GGUS:66865 [↗](#)) being investigated
- PIC
  - ◆ ATLAS data recovery progressing steadily, files being copied at a steady rate
    - ◇ concerns ~300 TB
    - ◇ 60% was back online on Friday
    - ◇ 40% is offline, but being copied behind the curtain
    - ◇ expected to finish in 2-3 days
    - ◇ final list of damaged files expected to be a few k
  - ◆ should we stop this and concentrate on the high-priority files instead?
    - ◇ not clear if that would be faster than 2-3 days!
    - ◇ Elena: we will discuss it and update the ticket
- RAL
  - ◆ LHCb data loss: ticket in progress, post-mortem will be produced
  - ◆ disk server reconfiguration for CMS OK (change of Puppet master node)
  - ◆ outage of LFC/FTS/3D database for Oracle upgrade on Wed Feb 9
- CASTOR
  - ◆ 2h downtime on Monday for EOS-ATLAS
- dashboards - ntr
- databases - ntr
- grid services - ntr

AOB:

- Stefan: ELOG crash reported by CMS yesterday has been looked into
  - ◆ the automatic restart by Lemon failed 3 times
  - ◆ the operator then started it manually according to procedure
  - ◆ the service is known to suffer crashes occasionally

## Friday

Attendance: local (AndreaV, Maarten, Stefan, Ewan, Mattia, Eva, Nicolo, Jan, Roberto); remote (Francesco, Jon, Xavier, Tore, Maria Francesca, Rolf, Gonzalo, Onno, Tiju, Suijian, Rob, Elena).

Experiments round table:

- ATLAS reports -
  - ◆ T0/CERN, Central services
    - ◇ srm-atlas.cern.ch had reduced availability ( 75%) for 20 min after midday today.
    - ◇ PVSS data replication has problem since this morning. ATLAS dba experts are investigating.
    - ◇ [AndreaV: saw an email thread about ADCR, what is the status? Eva: there is a bug for LOBs in Oracle 10.2.0.5, which has affected PANDA. Got a patch from Oracle and tested on an integration RAC that applying it does not cause extra problems (the PANDA issue could not be reproduced, so could not test if the patch fixes the problem. The patch will be applied on ADCR next Tuesday.)
  - ◆ T1s
    - ◇ PIC GGUS:66409 [↗](#): migration procecc is ongoing, some file transfers and jobs in pic cloud are affected.
    - ◇ FZK GGUS:66783 [↗](#): 15% of atlas jobs are failing because of network problem. Tests are set up between WNs in FZK and atlas server in CERN to reproduce the problem.

- ◇ BNL: reduced efficiency in file transfers because BNL\_SCRATCHDISK is almost full, many datasets are pending to be cleaned, BNL-OSG2\_SCRATCHDISK is put in first priority (within BNL) for the central deletion service. [Mattia; found problems with jobs from SAM tests, Alessandro is investigating. Maarten: this is now understood and solved, the CRLs were out of date].
- CMS reports -
  - ◆ CERN
    - ◇ No data taking
    - ◇ Ongoing debugging effort with IT-DB to understand DB connection issues in T0 workflow injectors
    - ◇ CASTORCMS DEFAULT occasionally degraded - expected due to recall from tape of 2M old tiny files which are going to be archived elsewhere SAV:105027 [↗](#)
  - ◆ T1s
    - ◇ MC reprocessing with 3\_9\_7 is in the tails.
  - ◆ T2s
    - ◇ 7TeV production ongoing.
    - ◇ T2\_UK\_SGrid\_Bristol - in SCHEDULED downtime for one week for machine room maint
    - ◇ T2\_DE\_RWTH: SAM CE tests: job submission failing SAV:119047 [↗](#)
- ALICE reports [↗](#) -
  - ◆ T0 site
    - ◇ CAF nodes: remaining nodes delivered
    - ◇ Migration to SLC5 of the xrootd servers, prepared and tested. Thanks a lot to Ignacio Reguero from IT-DSS. [Maarten: the migration is expected to be completed by the end of next week.]
  - ◆ T1 sites
    - ◇ Nothing to report
  - ◆ T2 sites
    - ◇ Nothing to report
- LHCb reports -
  - ◆ Experiment activities:
    - ◇ MC productions: due to an internal issue in DIRAC no pilot jobs have been submitted to pick up payloads. System almost emptied.
  - ◆ New GGUS (or RT) tickets:
    - ◇ T0: 0
    - ◇ T1: 0
    - ◇ T2: 0
  - ◆ Issues at the sites and services:
    - ◇ T0
      - Sent a request to relevant CASTOR people to help our Data Manager to understand/reconstruct why some files present in high level catalogs are not available physically on the Storages (not only at CERN). [Jan: please open a ticket, too.]
    - ◇ T1
      - RAL: after data loss, the correct re-replication is currently being checked by LHCb data manager (GGUS:66853 [↗](#))
      - SARA: yesterday user jobs failing. Increased the timeout in the dcap connection. (GGUS:66287 [↗](#))
      - GridKA: fraction user jobs failure rate a bit larger than usual. It seems a load issue. [Mattia: some SRM tests are failing, too.]
    - ◇ T2 site issues: :

· NTR

Sites / Services round table:

- CNAF: ntr
- FNAL: ntr
- KIT: ntr
- NDGF: will be unavailable for ATLAS between next Monday 7th 8am UTC and Tuesday 8th 8am UTC, due to site downtime in Oslo.]
- IN2P3: ntr
- PIC:
  - ◆ Concerning the ATLAS corruption, received the list of files in the GGUS ticket. These are 70k files, most are logfiles: should we handle logfiles at lowest priority? [Elena: yes logfiles are less important, please start by recovering the others.]
  - ◆ Next Tuesday 8th, short intervention (< 4h) to upgrade Oracle backend for FTS.
- NLT1: ntr
- RAL: next Mon intervention to upgrade GridFTP for CASTOR.
- ASGC: ntr
- OSG: ntr
  
- Services: ntr

AOB: none

-- JamieShiers - 26-Jan-2011

---

This topic: LCG > WLCGDailyMeetingsWeek110131

Topic revision: r16 - 2011-02-04 - AndreaValassi



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback