

# Table of Contents

<b>Week of 111031.....</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.....	6
Thursday.....	8
Friday.....	9

# Week of 111031

## 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(Yuri, Jamie, Massimo, Jan, Lola, Mattia, Steve, Jhen-Wei, Eva, Pepe, Maarten);remote(Onno, Lisa, Mette, Michael, John, Pavel, Giovanni, Rob, Burt).

Experiments round table:

- ATLAS reports -
- T0/Central services
  - ◆ ATLAS-AMI-CERN DB issues (Friday evening and Saturday). Number of sessions hit the limit of 400 and low availability in SLS. Seems fixed.
  - ◆ CERN PROD DATA DISK space token unavailable in SLS for short while. Transfer progressed ok so... Came back to normal Sunday afternoon so probably temp issue. SRM i/f for EOS? [ Jan- seems to be high load. ATLAS SLS test gave up after 15" trying to get space tokens. ]
- T1 sites
  - ◆ NDGF-T1-MCTAPE staging failures (Sat.): "No pool candidates available/configured/left for stage". GGUS:75815. - affects efficiency or Nordugrid cloud
  - ◆ CNAF transfer failures (Sat.): FTS connection timeout. GGUS:75816. 3 problematic channels overloaded with jobs. Two other errors - unable to connect to gridftp storm server; 2nd: invalid path.
  - ◆ TRIUMF: job failures with stage-in errors (Sat.) "File not online. Staging not allowed".

GGUS:75821 solved: issues with one of DDN controller, the storage expert tried to restarted at ~1:50 am on Sunday. This helped, test jobs succeeded and site is back to production.

- ◆ RAL-LCG2 massive T0 export and SAM SRM test failures. ALARM GGUS:75823 filed at 9pm on Sat. In ~1h. RAL team reported this caused by the high load. Reduced FTS file limits for ATLAS and the number of Atlas jobs allowed to run. Still job failures due to castor issues. Unscheduled DT announced on Sunday till Mon12:30. Looks like extended to 17:00. [ John - at the moment still in downtime; don't fully understand. In communication with ATLAS and CERN Castor team. ] [ Steve - T0 webserver ok this morning. If SRM e-p is down at RAL this might be problem. ]
  - ◆ Michael: ATLAS-wide 100K production jobs and 300K analysis jobs in holding. # finished jobs almost 0. Anything known? Yuri - believe it is ATLAS internal issue aware since 07:30. Experts investigating. Almost understood - config problem with Panda which probably was caused by config change by Lyon people. Trying to reset the status back to previous state. Helped but not completely resolved. Will take some time to recover.
  - T2 sites
    - ◆ ntr
- 

- CMS reports -
- LHC / CMS detector
  - ◆ The 2011 pp run is gone! Next stop HI runs...
- CERN / central services
  - ◆ CMSR Oracle instance spontaneous reboot problem, GGUS:74993, kept open to follow up with increased logging information, but no other spontaneous reboot seen since then...
- T0 and CAF:
  - ◆ cmst0 : very busy processing data !
  - ◆ On Sunday, all CE's were failing CMS SAM tests at T0\_CH\_CERN. The Error was "File was not found, could not be opened, or is corrupted". Originally thought it was a problem on CERN side, we fired an ALARM GGUS ticket (GGUS:75833). After some debugging it was learnt the SAM dataset was not at CERN anymore!! The SAM sample was deleted from CERN as part of the recent **RelVal** deletion campaign. Placed transfer request and closed ticket. Future SAM samples need to have better naming (not include RelVal on the name, to avoid these mistakes).
- T1 sites:
  - ◆ [T1\_FR\_CCIN2P3]: exports to T2s. GGUS:75397. The backlog is getting slowly reduced (91.5 TBs queued 1 week ago - 50 TBs atm).
  - ◆ [T1\_IT\_CNAF]: CREAM reporting dead jobs @ CNAF T1 as REALLY-RUNNING. GGUS:75648. JobIds provided and experts debugging the issue with the blah developers. It seems blah had missed to update the status of these jobs. Anyway it seems that all these jobs had been killed by the batch system because they reached the batch system memory limit (2.5GB)... Not much progress recently. [ Giovanni - waiting still for blah developers ]
  - ◆ [T1\_IT\_CNAF]: 1 corrupted file at CNAF. Savannah:124310. This is the only replica of the file. It might need an invalidation. DataOps taking a look to it.
  - ◆ [T1\_TW\_ASGC]: file access problem, GGUS:75377. In the end, there was a problematic file. We recovered it from PIC. Now we need to certify we can run happily on this file there.
  - ◆ [T1\_TW\_ASGC]: CMS Fall11 jobs have problems opening some files. GGUS:75784. There were problems staging the files. Now, that's been forced. After some work, there was a problem with one file, which was recovered from PIC. Now we need to certify we can run happily on this file there.
  - ◆ [T1\_TW\_ASGC]: transfer of the dataset /ZbbToLL\_M-30\_7TeV-mcatnlo-photos/Summer11-PU\_S4\_START42\_V11-v1/AODSIM from T1\_TW\_ASGC to some T2 sites. This transfer is stuck for a few days. GGUS:75780. Only 1 file is missing to land at 1 T2 site (out of 3, original on the transfer request).
  - ◆ [T1\_TW\_ASGC]: In particular, it seems we have problems exporting data in ASGC, as seen

in the Site Readiness as well

([http://lhweb.pic.es/cms/SiteReadinessReports/SiteReadinessReport.html#T1\\_TW\\_ASGC](http://lhweb.pic.es/cms/SiteReadinessReports/SiteReadinessReport.html#T1_TW_ASGC)).

We have additional GGUS tickets for this: GGUS:75834 and GGUS:75630. [ Jhen-Wei - as reported last week 1 d/s was broken, disabled it and tried to restage to other d/s. Fixed d/s yesterday. Since Saturday started to change old d/s and put offline to do some reconfig. Draining affected production so asked to stop. Stopped export to other sites this morning and just restarted now.

- Tier-2s:
    - ◆ Business as usual...
  - Others:
    - ◆ On Monday we spotted out a problem with downtimes tracing in the SSB. There is a bug which affected, at least CNAF, which was marked as being in Unscheduled Downtime for some days, which was not the case. Savannah:124240.
    - ◆ Migration from LCG-CE to CREAM not caught by nagios/dashboard. Savannah:124289 --> The 'latest result' page is ok, however the historical plots same CE for both flavours displayed there, if the same CE name was used for a transition LCG-CE<->CREAM-CE. To be fixed.
    - ◆ Bug affecting T1\_UK\_RAL, Savannah:124364 [↗](#), Decommissioning of a LCG-CE brought the site down for 2 months in SSB!
    - ◆ T2\_BR\_UERJ, Savannah:124025 [↗](#). The new outage was not visible in SSB... Collectors were stuck last night. Shouldn't CRC receive alert email for this? (Note: it seems it was not a normal stuck-event).
- 

- ALICE reports [↗](#) -
  - ◆ central services were affected by a power cut on Sat and another one Mon morning; (almost) all OK now - # running jobs has decreased; problem with cluster monitor that does not restart at sites - working on it.
- LHCb reports -

Sites / Services round table:

- NL-T1 - this w/e we had again stage requests that were stuck; had to restart dCache pools and hence some file transfers might have failed; in contact with dCache devs on this issue
- FNAL - ntr
- NDGF - ntr
- BNL - ntr
- RAL - in addition to ATLAS woes have network at risk tomorrow; very minimal
- KIT - public holiday tomorrow
- CNAF - nta
- ASGC - ntr
- OSG - ntr
  
- CERN storage: we are about to start tests with ALICE for CASTOR for HI. Detector will produce "fake data". end to end test. Should last 24h. Peak data rates higher than last year - want to be sure if LHC performs well can digest all data. Jan - SLS internal tests will get rerouted to experiment stagers instead of castor public which might cause a glitch.

AOB:

## Tuesday:

Attendance: local(David, Eva, Jan, Jhen-Wei, Maarten, Maria D, Mark, Nilo, Pablo, Steve);remote(Giovanni, Ian, Jeremy, John, Lisa, Mette, Michael).

Tuesday:

Experiments round table:

- ATLAS reports -
  - ◆ Since Sunday evening large number of jobs were staying at holding state in panda system. This problem was caused by misconfiguration for LYON. Experts fixed it Monday morning. But a lot of backlog to make load to DQ2. We offline all clouds Monday evening to wait jobs finish. Tuesday morning we online all clouds and panda system is back to production.
  - ◆ T0/Central services
    - ◇ ATLR DB was down this morning. Informed DBA and problem was fixed quickly. Thanks.
    - Eva: the storage hung after a normal rebalancing operation, cured by a restart
  - ◆ T1 sites
    - ◇ NDGF-T1-MCTAPE staging failures (Sat.): "No pool candidates available/configured/left for stage". GGUS:75815. The same problem on Tuesday.
      - Mette: it is being looked into, no news yet
    - ◇ CNAF transfer failures (Sat.): FTS connection timeout. GGUS:75816. 3 problematic channels overloaded with jobs. On Mon. other errors:"Unable to connect to gridftp-storm-atlas.cr.cnaf.infn.it" and "invalid path" On Tuesday other error appeared "[DESTINATION error during TRANSFER\_PREPARATION phase: [INVALID\_PATH] Invalid path]" so ticket was assigned. GGUS:75872 solved: The limit on the number of directory was reached. Now cleanup of empty directories is ongoing.
      - Maarten: it actually may be the responsibility of ATLAS to clean up such empty directories
    - ◇ RAL-LCG2 massive T0 export and SAM SRM test failures. ALARM GGUS:75823. Recovered CASTOR on Monday afternoon and ran without serious errors last night. Increased the atlas FTS channels here at RAL to 50% of their normal capacity. The batch farm limits have been raised to 1500 jobs for atlas production and analysis. Will monitor the situation during the day.
    - ◇ TAIWAN-LCG2 is set offline for LFC migration tomorrow (Nov. 2)
  - ◆ T2 sites
    - ◇ ntr
  - ◆ Other business
    - ◇ New project name 'data11\_hip' defined which can be included in Tape Family.
- CMS reports -
  - ◆ LHC / CMS detector
    - ◇ Magnet coming down detector maintenance
  - ◆ CERN / central services
    - ◇ CMSR Oracle instance spontaneous reboot problem, GGUS:74993, kept open to follow up with increased logging information, but no other spontaneous reboot seen since then...
  - ◆ T0 and CAF:
    - ◇ cmst0 : Last reco jobs should launch today
  - ◆ T1 sites:
    - ◇ [T1\_UK\_RAL] SAV:124394: Pilot Problems RAL Tier 1 Production
  - ◆ Tier-2s:
    - ◇ Business as usual...
  - ◆ Others:
    - ◇ Ian Fisk CRC
    - ◇ Maarten: any arrangements with CASTOR to prepare for HI run?
    - ◇ Ian: the HI data rate and volume will be rather normal this year thanks to zero suppression already happening online
    - ◇ Jan: nonetheless an e-mail to the CASTOR team would be appreciated whenever

some HI test activity is started, to correlate that with changes, if any, in the behavior of CASTOR

- ALICE reports [↗](#) -
  - ◆ yesterday's troubles with the Cluster Monitor VOBOX site services were cured by repairing an AliEn central service table that got corrupted due to yesterday's power cut
- LHCb reports -
  - ◆ Experiment activities
    - ◇ Prompt reconstruction and stripping at CERN and Tier-1 sites.
    - ◇ Calibration pushed through at CERN over the weekend
    - ◇ 1st round of reprocessing at T1 sites and T2 sites almost over (some tail going on - especially at GridKa)
    - ◇ Planned reprocessing at the end of this week when calibrations are complete.
  - ◆ T0
    - ◇ Done calibration runs now switched back to Stripping
    - ◇ Issue with possibly broken Raw Tape now resolved. See [GGUS:75859](#) [↗](#)
  - ◆ T1 sites:
    - ◇ Some inconsistencies with reported SE use and actual SE use (after counting files) at IN2P3. See [GGUS:75158](#) [↗](#)

Sites / Services round table:

- ASGC - ntr
- BNL
  - ◆ a 3-day maintenance will start on Mon Nov 7 at 07:30 local time lasting till ~14:00 on Wed:
    - ◇ network maintenance Mon-Tue will affect all computing services
    - ◇ USATLAS Oracle servers will be moved to a new data center
    - ◇ dCache: change from PNFS to Chimera name server
    - ◇ HW/SW upgrades of grid service nodes (CE, VOMS, ...)
- CNAF
  - ◆ ATLAS FTS tickets being worked on
  - ◆ CMS problem with CREAM-CE (see Monday report): still waiting for the BLAH developers response
- FNAL - ntr
- GridPP - ntr
- NDGF - ntr
- OSG - ntr
- RAL
  - ◆ FTS channels for ATLAS were increased to 75% capacity; if still OK tomorrow go to 100%
  - ◆ CMS Savannah ticket rather looks like a CMS issue
  - ◆ tape robot intervention went OK
- CASTOR/EOS:
  - ◆ ALICE HI data test: write+read overload yesterday evening (on stager + stager DB), now running OK (ongoing tuning by ALICE DAQ)
  - ◆ CASTOR stager databases for ATLAS, CMS, ALICE, CERNT3 got ORACLES patches. LHCb, central NS databases get them tomorrow.
  - ◆ CASTOR central nameserver transparent update to 2.1.11 tomorrow (see [GocDB](#))
  - ◆ EOSCMS update to 01.0-41 tomorrow at 10:00 (30 min intervention, got OK from CMS)
  - ◆ all CASTOR stager instances get transparent updates to 2.1.11-8 on Thu
  - ◆ SLS SRM-ATLAS, SRM-CMS, SRM-LHCB glitch yesterday was pure monitoring issue, SLS tests now run on correct CASTOR instance (and no longer on CASTORPUBLIC).
- dashboards

- ◆ looking into the recent tickets; some have already been closed
- ◆ working on 2 tickets related to downtimes, fixes expected this week
- ◆ 2 tickets related to the old SAM system have less priority in view of ongoing work on ACE that should take over in the near future
- databases
  - ◆ having validated the latest Oracle security patches we will deploy them in production, starting with the CMS online and production DBs on Thu
- GGUS/SNOW - ntr
- grid services
  - ◆ Monday's t0export FTS web service upgrade led to a problem for 1 user, now resolved

AOB:

## Wednesday

Attendance: local(Mark Slater, Maria, Jamie, Maarten, Jhen-Wei, Jan, David, Alessandro, Steve, Luca, MariaDZ, Ian);remote(Michael, Mette, John, Giovanni, Rolf, Pavel, Rob, Ron, Lisa).

Experiments round table:

- ATLAS reports -
  - T0
    - ◆ CERN-PROD has CASTORATLAS headnode hardware intervention this afternoon (Wed, 2 November, 13:00 – 14:00 UTC). Failover required.
  - T1 sites
    - ◆ NDGF-T1-MCTAPE staging failures (Sat.): "No pool candidates available/configured/left for stage". GGUS:75815. Under investigation. [ Mette - we think problem is solved; seems that a pool lost its config after restart and didn't notice as wasn't in Nagios. Think ok now; people just taking a look now. Will update GGUS when we have final confirmation ]
    - ◆ INFN-T1 transfer failures (Sat.): FTS connection timeout. GGUS:75816. 3 problematic channels overloaded with jobs. On Mon. other errors:"Unable to connect to gridftp-storm-atlas.cr.cnaf.infn.it" and "invalid path". The limit on the number of directory was reached. Now cleanup of empty directories is ongoing.
    - ◆ RAL-LCG2 massive T0 export and SAM SRM test failures. ALARM GGUS:75823 . The Atlas FTS channels and batch farm limits have been returned to normal here at RAL. Unless there are further developments, I propose solving this ticket this afternoon.
    - ◆ TAIWAN-LCG2 scheduled downtime until 11:00am UTC for power construction. Set another 4-hour unscheduled downtime (11:10 - 15:10) to complete all.
    - ◆ TAIWAN-LCG2 is set offline whole day for LFC migration.
  - T2 sites
    - ◆ ntr
- 

- CMS reports -
- LHC / CMS detector
  - ◆ Quiet
- CERN / central services
  - ◆ We had an SLS failure yesterday that marked but Castor Default and SRM as bad, though both appeared to be working in practice. We submitted a team ticket and both recovered. [ Jan - due to short timeouts in probes. Routing srm tests through castor default. CMS invited to propose better pool to use for tests ]
  - ◆ One of the PhEDEx Data service APIs is broken after yesterday's update. Under investigation
- T0 and CAF:
  - ◆ cmst0 : Started investigating some Tier-0 backfill workflows.

- T1 sites:
    - ◆ [T1\_FR\_IN2P3] Savannah ticket #124418: Bad files at IN2P3. Updated with GGUS bridge
      - ◇ Unavailable data files - 4 files only at IN2P3.
    - ◆ [T1\_IT\_CNAF] sr #124417: No new CMS Tier 1 Production jobs submitted at CNAF
      - ◇ Looks like a CMS factory issue, no site action required at this time
  - Tier-2s:
    - ◆ Business as usual...
  - Others:
    - ◆ Working to close old Savannah tickets
- 

- ALICE reports [↗](#) -
    - ◆ NTR
- 

- LHCb reports -
  - Experiment activities
    - ◆ Prompt reconstruction and stripping at CERN and Tier-1 sites.
    - ◆ Last few calibration jobs going through today
    - ◆ Last few hundred jobs of 1st reprocessing running now
    - ◆ Looking to start 2nd stage reprocessing on Friday (maybe Monday or maybe even tomorrow!)
  - T1 sites:
    - ◆ Update on site inconsistencies: IN2P3 seems to have the only significant problem. Other T1s seem OK. See [https://ggus.eu/ws/ticket\\_info.php?ticket=75158](https://ggus.eu/ws/ticket_info.php?ticket=75158)
    - ◆ We are exceeding TAPE pledges at some sites and so, as a short-medium term solution we are going to ban ARCHIVE at sites where this is a problem.
    - ◆ We are still having occasional issues with respect to staging and job failures at various sites with input data resolution which we're looking into.
    - ◆ GridKA - gap in SAM tests for GridKA, Stefan looking into it, might be heavy load and tests not running
- 

Sites / Services round table:

- BNL - ntr
- NDGF - ntr
- RAL - ntr
- CNAF - concerning problem of CMS with CREAM CE; blah problem.
- IN2P3 - we had a major problem with our tape robot on Monday evening; still has some consequences. Some tapes still unavailable, still working on it. Suppose CMS problem might be related to this. ALICE jobs - ongoing problem; they are queuing up and we are working on it
- KIT - ntr
- NL-T1 - ntr
- FNAL - ntr
- OSG - ntr
  
- CERN FTS T0 export service: Proposal to update the FTS agent nodes for **fts22-t0-export.cern.ch** on **Wednesday 9th November from 10:00**. The update is from SLC4 -> SLC5 and to FTS version **3.2.1-2.sl5**. This is the FTS version that has been running on the tier2 service for many months. This is final part of the SLC4 -> SLC5 migration for FTS. During the intervention of a couple of hours **the Webservice will accept new jobs but they will queue up**. To be confirmed at the end of this week. - feedback suggests that Monday would be better
  
- CERN storage: had short update on EOS CMS and transparent updates on all CASTOR instances tomorrow.



AOB:

- GGUS tickets for tomorrow's T1SCM. Any input?

## Thursday

Attendance: local(Jan, David, Andrea, Steve, Maria, Jamie, Zbyszek, Ian, Jhen-Wei, Mark, Lola, MariaDZ);remote(Michael, Mette, Lisa, Ronald, Gareth, Rolf, Ronald).

Experiments round table:

- ATLAS reports -
- T0
  - ◆ High CPU load in ADCR DB observed at 4pm. Contacted DBA and problem was fixed within one hour. Thanks
- T1 sites
  - ◆ NDGF-T1-MCTAPE staging failures (Sat.): "No pool candidates available/configured/left for stage". GGUS:75815. No more error. Can ticket be closed?
  - ◆ INFN-T1 transfer failures (Sat.): FTS connection timeout. GGUS:75816. 3 problematic channels overloaded with jobs. On Mon. other errors:"Unable to connect to gridftp-storm-atlas.cr.cnaf.infn.it" and "invalid path". The limit on the number of directory was reached. Is cleanup of empty directories still ongoing? No more error observed.
  - ◆ TAIWAN-LCG2 unscheduled downtime 3 Nov. 12:00 UTC to recover services from power maintenance.
  - ◆ TAIWAN-LCG2 is set offline whole day for LFC migration.
- T2 sites
  - ◆ ntr

- CMS reports -
- LHC / CMS detector
  - ◆ Quiet
- CERN / central services
  - ◆ Oracle update today was successful
- T0 and CAF:
  - ◆ cmst0 : Working on Heavy Ion replays
- T1 sites:
  - ◆ [T1\_FR\_IN2P3] Savannah ticket #124418: Bad files at IN2P3. Updated with GGUS bridge [ Rolf - work on-going. The tape system problem started Monday evening and all robots went down suddenly. During intervention technicians were obliged to take several tapes out of the robot and a full update of the location database for the tapes needed to be done. Those tapes in bad locations had to be re-entered. This created a lot of trouble with HPSS / dCache which expected data to be on known tapes. Should be over since yesterday noon. If you still observe something problem is perhaps elsewhere. Don't have access to tickets for the moment]
    - ◇ Unavailable data files
  - ◆ [T1\_FR\_IN2P3] another report of unavailable transfers.
- Tier-2:
  - ◆ Bussiness as usual...
- Others:
  - ◆ Working to close old Savannah tickets

- ALICE reports [↗](#) -
  - ◆ NTR

- LHCb reports -
  - Experiment activities
    - ◆ Almost all reconstruction and reprocessing jobs now complete
    - ◆ We're trying to push the stripping through (large backlog at CERN)
    - ◆ New conditions DB tests running today so will hopefully start the next round of reprocessing tonight or tomorrow morning
    - ◆ CERN will be kept out of this to try to get the stripping complete
  - T0
    - ◆ FTS Channels both to and from CERN appear to have been set offline. Raised a ticket: [https://ggus.eu/ws/ticket\\_info.php?ticket=75936](https://ggus.eu/ws/ticket_info.php?ticket=75936) [ Steven - nearly solved, it is a monitoring issue. ]
- 

Sites / Services round table:

- NDGF - ntr
- RAL - we had a bit of a problem on batch system for one hour - problems with access to /tmp 11 - 12.
- IN2P3 - nta
- NL-T1 - ntr
- FNAL - ntr
  
- CERN Grid intervention is now **Tuesday** - on CERN SSB
  
- CERN storage - update on CASTOR public done today; in progress of doing one for CASTOR LHCb. Error in announcing CASTOR ATLAS so resched to 7-Nov
  
- CERN DB - as mentioned by CMS patched all CMS prod DBs. For other experiments scheduled Mon-Tue next week.

AOB:

## Friday

Attendance: local(Massimo, JhenWei, Ian, Zbigniew, David, Ueda); remote(Michael, Alessandro, Gonzalo, Lisa, Rob, Joel, Onno, Gareth, Mette, Rolf, Xavier).

Experiments round table:

- ATLAS reports -
  - ◆ T0
    - ◇ CERN-PROD transfer failures: "failed to contact on remote SRM [<http://srm-eosatlas.cern.ch:8443/srm/v2/server>]. Givin' up after 3 tries". GGUS:75953. Solved: The SRM was refusing to accept certificates issued by CERN. It is not yet understood what is the reason behind it. Might be a corrupted revocation list for the CERN CA. The problem is fixed now.
    - ◇ ATLAS databases rolling interventions for Oracle security updates 7,8,9 November. Will LCGR perform same action?
  - ◆ T1 sites
    - ◇ TAIWAN-LCG2 is set online
    - ◇ BNL downtime 7-9 November. 1) T0 export will be stopped today to avoid mis-operation during the weekend. 2) Set BNL offline in Panda, schedule under planing 3) DDM transfer will be excluded automatically according to downtime.
    - ◇ BNL also takes actions on FTS, batch queues accordingly.
  - ◆ T2 sites
    - ◇ ntr

- CMS reports -
  - ◆ LHC / CMS detector
    - ◇ Quiet
  - ◆ CERN / central services
    - ◇ Small cmsweb intervention today to deal with broken functionality
  - ◆ T0 and CAF:
    - ◇ cmst0 : First successful HI replays
  - ◆ T1 sites:
    - ◇ [T1\_TW\_ASGC] currently has a reasonable level of Job Robot aborts.
    - ◇ [T1\_FR\_IN2P3] and [T1\_DE\_KIT] were contacted by T2\_US\_Nebraska. Nebraska has a new PerfSonar system and is finding trouble.
  - ◆ Tier-2s:
    - ◇ Transfer errors were reported to T2\_US\_Vanderbilt, which is important to the HI run
  - ◆ Others:
    - ◇ Working to close old Savannah tickets
  
- ALICE reports [↗](#) -
  - ◆ NTR
  
- LHCb reports -
  - ◆ Experiment activities
    - ◇ Started tests for new reprocessing with 372 files last night
    - ◇ Mostly OK, but some failures at IN2P3 due to Conditions DB not being installed on AFS and not all the cluster being on CVMFS
    - ◇ Reprocessing is now being ramped up
    - ◇ Stripping is continuing at CERN: ~15000 to go at a rate of 3.5K per day
  - ◆ T0
    - ◇ There still seems to be FTS problems at CERN. The last few days there have been no successful transfers to/from CERN. Could be related to an authorisation error we're seeing - see ticket: [https://ggus.eu/ws/ticket\\_info.php?ticket=75936](https://ggus.eu/ws/ticket_info.php?ticket=75936). (Note that I reported this incorrectly yesterday as I didn't realise the plots I was looking didn't show failed attempts, only failed transfers!). Problem is fixed but can we have an incident report, please ?
  - ◆ T1
    - ◇ IN2P3 :We can't stage a few files at IN2P3 - Ticket opened: [https://ggus.eu/ws/ticket\\_info.php?ticket=75991](https://ggus.eu/ws/ticket_info.php?ticket=75991) (why after 5 hours it has not been acknowledged???)
    - ◇ GRIDKA : lost 2 RAW files (GGUS:75922)

Sites / Services round table:

- ASGC: CASTOR 2.1.11-6 will be installed next week (Mon-Tue).The CMS JobRobot errors have been tracked down to some misconfigured worker nodes.
- BNL: To avoid loss of capacity (US Tier2) during the BNL downtime cross-cloud capability of Panda will be used
- CNAF: Network problems this morning (DNS reconfiguration) now solved. ATLAS data transfer from PIC to CNAF affecting ATLAS: it seems due to overload of requests (investigating)
- FNAL: NTR
- IN2P3:Late response to tickets due to some GGUS tickets lost when entering the local tracking system
- KIT:Ask experiments to publish GGUS ticket number in the report. An incident report will be published (GGUS:75922 [↗](#))
- NDG: NTR
- NLT1:NTR

- PIC:NTR
- RAL:CASTOR JobMgr problem this morning (requets not fulfilled ~45' this morning around 8:00 UTC). Monday the BDII values for CASTOR will be fixed (due to a glitch RAL is publishing higher-than-available quantities).
- OS: NTR
  
- CASTOR/EOS:Monday CASTOR upgrade to 2.1.11-8 (ATLAS and CMS). Already published. Transparent
- Dashboards: NTR
- Databases: Next week is "patching week". Many interventions visible on the ITSSB. ATLAS mentioned that in case of an overrun of the intervention on the VOMS DB we could have the situation of a blackout due to the BNL intervention (DB upgrade in the morning, BNL going in maintenance in the afternoon).

AOB:

-- JamieShiers - 14-Sep-2011

---

This topic: LCG > WLCGDailyMeetingsWeek111031

Topic revision: r23 - 2011-11-04 - MaartenLitmaath



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