

# Table of Contents

<b>Week of 101108.....</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.....	9
Friday.....	11

# Week of 101108

## 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(Jerka, Maria, Jamie, Hurng, Peter, Maarten, Simone, Andrea, Luca, Dirk, Harry, Ignacio, Nilo, Ricardo, Alessandro, Eddie, Lola, Roberto, Massimo, MariaDZ);remote(Michael, Xavier, Jon, Alexander Verkooijen, Daniele, Gonzalo, Federico, Rolf, John, Rob).

Experiments round table:

- ATLAS reports -
  - ◆ SARA-MATRIX had BDII issues & no ATLAS jobs were running. GGUS ALARM ticket #63994 sent on Saturday evening. Ticket was acknowledged after 17 hrs. Bunch of emails sent to SARA contact email (the one in GGUS) and to site administrators along with GGUS. Can SARA please make sure that ALARM tickets are acknowledged a bit sooner? Two separate issues: BDII, and issue with local batchsystem. SARA cloud excluded from ATLAS DDM for night, cloud offline for production until Sunday noon. Issue fixed Monday morning. Thank you!
  - ◆ BNL: sunday morning lots of "dCache" errors in DDM. GGUS ALARM ticket #63999 sent. Issue fixed within several hours from report. Thank you!
  - ◆ INFN-T1: Saturday evening SRM failures. GGUS TEAM ticket #63995 (top priority) on Saturday evening. No response to the TEAM ticket until Sunday morning, we escalated to GGUS ALARM ticket #64000. Issue with storm - underlying GPFS cluster. Issue fixed within couple hours. Thank you! [ Daniele - INFN T1 StoRM issue. Didn't answer TEAM

ticket until Sunday morning as opened after midnight on Saturday. Problem was not with StoRM itself but underlying GPFS cluster that had problems due to human error. Actions taken to avoid such problems in future. ]

- ◆ Taiwan-LCG2: SRM issues, GGUS team ticket #63983
  - ◆ Lyon: No major problems over the weekend. Can you please comment on previous issues evolution? Thank you! [ Rolf - the jobs of ATLAS ran well this w/e as ATLAS increased timeout values. This was sufficient to get jobs through. Obviously this does not solve underlying problem - working on this. Investigation and tests on basis of LHCb and ATLAS jobs (LHCb have same problem). No more details for now but our Tier1 representative will give more complete report at next T1SCM. Any news will be reported here. ]
- CMS reports -
    - ◆ Experiment activity
      - ◇ First HI collisions over the week-end
    - ◆ CERN and Tier0
      - ◇ CMS reconstructed first HI events/tracks
      - ◇ Tier-0 : small SW worries for some Prompt Reco/AICa workflows at Tier-0, but patch was made and is now in place
      - ◇ SLS glitch this morning 10:50-11:20, but got resolved
    - ◆ Tier1 issues and plans
      - ◇ IN2P3-CC : Monte-Carlo Pile-up production now complete. Will immediately switch to Heavy Ion simulation.
      - ◇ More generally : processing data re-reco + skimming of new HI data at all Tier-1s
    - ◆ Tier2 Issues
      - ◇ T2\_RU\_RRC\_KI : CMS SAM-CE instabilities since several days, see <http://dashb-ssb.cern.ch/dashboard/request.py/siteviewhistory?columnid=4&time=168&dateF>, and still no solution/response by site admin (Savannah:117543 [↗](#))
    - ◆ AOB
      - ◇ Upgrade of CMSWEB cluster last Thursday Nov 4, in particular new rule of `http==>https` re-direction, cause problems of some CMS WebTools clients, eventually causing excessive memory usage (blockreplicas calls) of CMSR node hosting PhEDEx, that needed to be rebooted Friday Nov 5 at 16:38 CET.
        - Immediate action was to patch the most active clients (ProdAgent, CRAB already had an up-to-date version but was not deployed everywhere)
        - CMS Operators in alert state over the whole week-end to monitor the situation
        - will discuss this afternoon internally on short/long-term solution
      - ◇ Next CMS Computing Run Coordinator, reporting here from tomorrow on : Oliver Gutsche
  - ALICE reports [↗](#) - GENERAL INFORMATION: right now there are stable Pb beams, data taking is ongoing. It has to be followed up immediately by MC for the RAW data runs with the same detector conditions data, we are keeping the resources on hot standby. Some problems with AliROOT solved on the spot.
    - ◆ T0 - ntr
    - ◆ T1 - ntr
    - ◆ T2 - usual operations
  - LHCb reports - MC campaign all the weekend
    - ◆ T0
      - ◇ Transparent intervention on CASTOR (10-12 UTC) Maybe not really transparent? GGUS:64043 [↗](#) [ Ignacio - don't see anything in CASTOR, could be SRM, will look at it. Roberto - looks like a disk server problem. Massimo - one thing that looks a bit unusual at least 4 users reading data with more streams than LHCbPROD. Unusual

but could be unrelated. ]

- ◆ T1 site issues:

Sites / Services round table:

- RAL - SIR for problems for LHCb at the RAL Tier1. These appeared first as bad TURLs being returned, then an interruption to the service. SIR [↗](#)
- KIT - during w/e one NFS server died. High impact on other services, e.g. WLM service died also. Fixed today. Probably will have unstable WNs rest of week but will have an eye on it.
- FNAL - had a lot of trouble this morning with Alcatel. Currently asking for a pw! Today is our downtime, started 3 hours ago, continuing until at least 18:00 Chicago time (10 more hours). Availability problem of last week partially fixed. Some done LCG side not yet propagated to CMS. We detector that FNAL-KIT network links down Saturday for some time - no more info.
- BNL - Channel to CERN didn't work!) In addition we encounter problems with SAM BDII. Our probe last night detected that LCG BDII dropped all resources associated with BNL T1 around 09:00 pm. Recovered shortly after. Also encountered issues with SAM BDII. Unclear why. Opened OSG ticket 9526. Maarten - OSG tickets not automatically forwarded to right GGUS category - should be assigned to CERN. Rob - submitted ticket a couple of hours ago to GGUS:64039 [↗](#).
- NL-T1 - 3 things: in night Fri-Sat NIKHEF BDII had vm problems and had to be restarted. On Sat SARA BDII restarted. Our batch system had some problems - several thousand ALICE jobs submitted. Simone - problem to batch and BDII non-correlated? A: yes, not correlated.
- CNAF - nothing to add to ATLAS report
- PIC - 1/2 hour ago cooling incident in part of the data centre - had to switch off part of farm. Affecting ~400 jobs. More later..
- IN2P3 - nta
- RAL - quite a few things. Gareth mailed in SIR - see above. Problems over w/e with site BDII - locked up and failed some SAM tests. ATLAS disk server failed this am. Fabric looking at it. ATLASDATADISK out of service. An outage for LHCb on Wed - to upgrade d/s & o/s to 64bit. Outage to CASTOR CMS next week.
- OSG - ATLAS alarm this w/e. Several people woken up early Sunday. Don't mind responding if really alarms. Made sure ticket forwarded to BNL. Would like to find out if this was a real alarm issue. Michael - no doubt it was an alarm situation and needed to be fixed asap which it was. Jerka -
- CERN storage - CASTOR LHCb upgraded to 2.1.9-9 - nothing to do with problems seen since.

AOB:

- LHC news: *"Geneva, 8 November 2010. Four days is all it took for the LHC operations team at CERN to complete the transition from protons to lead ions in the LHC. After extracting the final proton beam of 2010 on 4 November, commissioning the lead-ion beam was underway by early afternoon. First collisions were recorded at 00:30 CET on 7 November, and stable running conditions marked the start of physics with heavy ions at 11:20 CET today."*

## Tuesday:

Attendance: local(Eddie, Roberto, Oliver, Lola, Andrea, Jamie, Huang, Luca, Harry, Ricardo, Alessandro, Miguel);remote(Rolf, Michael, Jon, Kyle, Daniele, Ronald, John, Christian Sottrup (NDGF) ).

Experiments round table:

- ATLAS reports -
  - ◆ Nov 8th - 9th (Mon, Tue)
    - ◇ last night, heavy ion collisions with stable beams, ATLAS in data taking mode.

Tuesday:

- ◊ p-p data reprocessing going on at T1s. It's about to submit 3rd batch of jobs to process the data taken during October (~217 TB in total)
- ◆ T0:
  - ◊ this night around 4:07 - 4:22 am, we observed many acronjob failures, most probably due to AFS problems. Has CERN observed something? [ Harry. Nothing in morning meeting. Alessandro - have observed also yesterday short glitches, last week too. ]
- ◆ T1s:
  - ◊ short SE glitches at INFN-T1 last evening (~6:00 pm CERN time). After the GGUS:64056 [↗](#) opened, problem seemed to be gone. [ Daniele - GGUS ticket solved. In case it comes again please let us know. ]
  - ◊ data reprocessing backlog at SARA due to the CE issue reported yesterday. Although the problem was fixed, ATLAS didn't reach reasonable amount of CPU share to catch it up. SARA reconfigured the share so now more jobs are able to run at SARA (thanks for the help). Still watching the progress and preparing alternative plans in case we need it for next batch of data reprocessing.
  - ◊ FZK increased the DATADISK space token, 100% ESD replication resumed.
- ◆ T2s:
  - ◊ copy-on-demand mechanism for ESD/DESD analysis implemented for UK cloud; therefore auto-replication ESD/DESD stopped
  - ◊ keep T2-T2 functional test to CA-ALBERTA-WESTGRID-T2 and CA-VICTORIA-WESTGRID-G2 for debugging the FTS + gridftp door setup that is causing T2-T2 transfer failures in CA cloud.
- CMS reports -
  - ◆ Experiment activity
    - ◊ HI data taking
  - ◆ CERN and Tier0
    - ◊ Tier-0 : Software configuration issues, being resolved by various patch releases [ higher load on castor but seems to perform well. ]
  - ◆ Tier1 issues and plans
    - ◊ re-processing pp data + skimming at all Tier-1s
    - ◊ received new PileUp re-digi/re-reco requests which will start soon at the sites
    - ◊ T1\_US\_FNAL: SE type srmv2 not published  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=64084](https://gus.fzk.de/ws/ticket_info.php?ticket=64084) [↗](#) - already talked to admins who are working on it.
    - ◊ T1\_DE\_KIT: implementation of new HI production roles:  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=64069](https://gus.fzk.de/ws/ticket_info.php?ticket=64069) [↗](#)
    - ◊ T1\_UK\_RAL: compare checksum of file with database value, file corrupt at T2 site:  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=64069](https://gus.fzk.de/ws/ticket_info.php?ticket=64069) [↗](#)
    - ◊ T1\_FR\_CCIN2P3: transfer problems to MIT, solved?  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=63826](https://gus.fzk.de/ws/ticket_info.php?ticket=63826) [↗](#)
  - ◆ Tier2 Issues [ Rolf - just to signal that ticket # in CMS report is not correct! corrected above - to be checked offline. Ticket re-opened by local CMS support - will follow up ]
    - ◊ T2\_RU\_RRC\_KI : CMS SAM-CE instabilities since several days, see <http://dashb-ssb.cern.ch/dashboard/request.py/siteviewhistory?columnid=4&time=168&dateF>, and still no solution/response by site admin (Savannah:117543 [↗](#))
    - ◊ T2\_RU\_JINR: Site invisible in BDII, SAM tests fail, Phedex agents are down,  
<https://savannah.cern.ch/support/index.php?117742> [↗](#)
    - ◊ T2\_FI\_HIP: Wrong gstat and gridmap information for the Finnish CMS Tier-2:  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=63956](https://gus.fzk.de/ws/ticket_info.php?ticket=63956) [↗](#)
  - ◆ Infrastructure
    - ◊ two savannah tickets (117700,117701) opened several (8) ggus tickets?
- ALICE reports [↗](#) - GENERAL INFORMATION: The amount of RAW data is still at the level that we

can reconstruct all runs Pass1@T0 quasi online and the data replication to T1s can follow the data taking (no backlog). Analysis is ongoing. The MC production has returned to normal.

- ◆ T0 site
  - ◇ Nothing to report
- ◆ T1 sites
  - ◇ FZK: The disk at SE was full, so it could not store more files. They are working on getting new storage capacity up, so issue will be probably solved by today
- ◆ T2 sites
  - ◇ Usual operations
- LHCb reports - MC campaign still going on, users jobs, reconstruction tail
  - ◆ T0
    - ◇ GGUS:64043 [GGUS:64043](#) (opened yesterday). Files could not be accessed. CASTOR support now asking to close it, we still see some failures. Re-checking on our side. [ Miguel - you are seeing timeouts right? In effect the # files being opened and their location on servers is causing users to block each other out. Users see some timeouts - maximum # of connections reached. Roberto - user addressing huge number of files? A: that is a special case - user activity conflicting with each other on default pool. Would be better if user activity could be moved to xroot. Other option is to add some more h/w. This would have more limitations than first option. Roberto - do not have control on users doing analysis using lxbatch. Will encourage usage of xroot. ]
    - ◇ GGUS:63933 [GGUS:63933](#) (opened 5 days ago). Shared area slowness causing jobs to fail. Seems to be understood. Can probably close - just wait to see if there are some more problems in this area.
  - ◆ T1 site issues:
    - ◇ NTR
    - ◇ [ Eddie - still tests timing out with shared area of IN2P3. Both CE and SRM. Still many errors from many tests. Rolf - shared area problem - a lot of things are ongoing there. Have isolated some WNs to see if only LHCb is working on them. Try to simulate heavy load for LHCb setup. On concrete issue mentioned don't know if this has been seen here. Ticket reference? Roberto - no ticket opened on recent SAM failures. Rolf - report at next T1SCM and also MB. ]

Sites / Services round table:

- IN2P3 - nta
- BNL - all T1 services running smoothly. Problems with SAM BDII continue. Meanwhile found out that there are 2 used in round-robin mode. 1 fails constantly as it has no info on BNL resources in particular SE. Communication Hiro-Ale. Ale opened another ticket. Going through OSG and Rob had already filed a GGUS ticket. Close one ticket or cross-reference? Ale- will close it.
- FNAL - had our downtime yesterday, generally successful. A few loose ends, e.g. BDII published as Oliver mentioned. Reported to me that FNAL-KIT circuit restored yesterday. Faulty line card in AMS on US-LHCNET. FNAL runs a central PhEDEx service for US T3s. Node that does that crashed and is broken. All circuits to US T3s down. Replacing node.
- CNAF - there will be power intervention on Friday in the morning. Should not affect farm as diesel generators. Set an AT RISK but not more.
- NL-T1 - ntr
- RAL - 3 things today. Issue with ATLAS disk server still going on. Local ATLAS contact in touch. 1 m/c that serves site BDII out of DNS alias so only 1 left. Scheduled downtime tomorrow for LHCb SRM to upgrade to 64 bit.
- NDGF - upgrading dCache. Going according to plan will finish tomorrow around noon. During this time some data may be unavailable.
- OSG - noticed a ticket exchange problem with FNAL which we are currently trying to diagnose. Any tickets to US CMS might not arrive immediately. (Good says Jon!)

- CERN Grid services: LSF - need to schedule an intervention to replace h/w of 2 masters. Before end year? End Jan? Please suggest timeframe. Preferably beginning December if possible. acronjob failures - ticket svp. SAM BDII failures - please update ticket with info that only 1 of BDII is bad.
- CERN DB - replication of CMS online to offline stopped 07:15 - 09:15 due to problem with maintenance job - restarted manually. Investigating.
- CERN Dashboards - for ATLAS DDM dashboard scheduled downtime today at 10:00 for 25' All went smoothly - new version has statistics on data transfer rates.

AOB:

- Next T1SCM this coming Thursday, day after tomorrow.

## Wednesday

Attendance: local(Huang, Maria, Jamie, Andrea, Alessandro, Luca, Lola, Eddie, MariaDZ, Miguel, Ricardo, Carlos);remote(Xavier, Oliver, Onno, Michael, Kyle, Joel, Gang, Rolf, John, Christian).

Experiments round table:

- ATLAS reports -
  - ◆ data taking on 17x17 bunches beam last night. data taking going on.
  - ◆ T0:
    - ◇ 3rd instance of offline (ATLR) database highly loaded. Panda system is the main user. atlas-dba and Panda experts notified. The load was gone this morning.
    - ◇ CERN HOTDISK saturated (GGUS:64095 [↗](#)): [ working on adding a few extra disk servers to the pool. A lot of access to the exact same file. A limited amount of I/O operations and re-reading the same file... Advice: can ATLAS for future find a more scalable way to distribute this 1-2 files needed by all jobs? Ale - long standing issue. ATLAS are working on this. HOTDISK setup for just this purpose. More problems seen at CERN than at T1s. When could those disk servers be available? Miguel - identified servers; draining. May take a few hours - today or tomorrow. ]
      1. hitting network limitation for the whole pool. adding more machines to the pool?
      2. c2atlassrv201 had run out of /var causing xrootd redirector to fail. problem fixed after cleaning it up.
    - ◇ LSF master node hardware upgrade mentioned yesterday: intervention duration?
  - ◆ T1s:
    - ◇ INFN-T1: short LFC glitch this morning.
    - ◇ RAL: disk corruption on one of the disk server causing data lost. Site is trying to recover the files only available at RAL. Any progress? [ Brian - have recovered some of the files from disk server in question and repopulating back into CASTOR. An additional list for which we could also try to recover. But if ATLAS want to declare lost sooner - another 8K files would be declared lost. Ale- ok if those files are declared lost ATLAS can handle the situation. Fine if you declare that list lost today or else wait until tomorrow morning ]
    - ◇ SARA reprocessing backlog: job finish rate is doubled after reconfiguring the share and utilizing NIKHEF CPUs.
    - ◇ IN2P3: couple of data transfer failures collected at GGUS:64123 [↗](#)
- CMS reports -
  - ◆ Experiment activity
    - ◇ HI data taking, impressive ramp-up

- ◆ CERN and Tier0
  - ◇ Tier-0 : Software configuration issues, being resolved by various patch releases
  - ◇ Castor: putting high peak load on system, workflows under investigation to prepare for high luminosity running
- ◆ Tier1 issues and plans
  - ◇ re-processing pp data + skimming at all Tier-1s
  - ◇ received new PileUp re-digi/re-reco requests which will start soon at the sites
  - ◇ T1\_US\_FNAL: SE type srm2 not published  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=64084](https://gus.fzk.de/ws/ticket_info.php?ticket=64084) -> solved
  - ◇ T1\_DE\_KIT: implementation of new HI production roles:  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=64069](https://gus.fzk.de/ws/ticket_info.php?ticket=64069) -> still open
  - ◇ T1\_UK\_RAL: compare checksum of file with database value, file corrupt at T2 site:  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=64021](https://gus.fzk.de/ws/ticket_info.php?ticket=64021) -> ticket was bridged unnecessarily to GGUS, GGUS ticket closed
  - ◇ T1\_FR\_CCIN2P3: transfer problems to MIT, solved?  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=63826](https://gus.fzk.de/ws/ticket_info.php?ticket=63826) -> contacted local site contacts and waiting for feedback, last reply to ticket from 8th Nov: "Concerning transfers exporting from IN2P3 to other sites ( like MIT who opened this ticket), We still see the same errors : AsyncWait, Pinning failed"
- ◆ Tier2 Issues
  - ◇ T2\_RU\_RRC\_KI : CMS SAM-CE instabilities since several days and still no solution/response by site admin  
<https://savannah.cern.ch/support/index.php?117543>
  - ◇ T2\_RU\_JINR: Site invisible in BDII, SAM tests fail, Phedex agents are down,  
<https://savannah.cern.ch/support/index.php?117742> -> fixed and closed
  - ◇ T2\_FI\_HIP: Wrong gstat and gridmap information for the Finnish CMS Tier-2:  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=63956](https://gus.fzk.de/ws/ticket_info.php?ticket=63956), more explanation:
    - T2 FIN resources appear in BDII as part of NDGF-T1 because we need to translate the ARC information system to Glueschema and into BDII format and this is done by NDGF
    - Site has an independent CMS dCache setup which is not part of the distributed NDGF T1 dCache setup so that cannot be published by NDGF, but is instead published by a CSC BDII server.
    - CMS applications and WLCG accounting works fine with this setup, but the WLCG monitoring is reporting non CMS resources for the T2 Fin site which obviously is wrong.
    - The ticket is about having WLCG report the right resources for us. The short term solution is to for WLCG to change configuration settings to get this corrected.
  - ◇ T2\_US\_Caltech failed SAM tests  
<https://savannah.cern.ch/support/index.php?117760>, traced back to "SE had reached 100% capacity due to a heat-induced controlled shutdown of some storage nodes, as well as a large storage node that crashed last night" -> fixed and closed
  - ◇ T2\_FR\_IPHC failed SAM tests  
<https://savannah.cern.ch/support/index.php?117769>, problem with the SE\_mysql server -> fixed and closed.
- ◆ Infrastructure
  - ◇ two savannah tickets (117700,117701) opened several (8) ggus tickets? -> bridging was not necessary, GGUS tickets closed, multiplication of tickets will be followed up offline
- ALICE reports - GENERAL INFORMATION: Pass1@T0 and data replication to T1s are performing quite well. Analysis trains and three MC production cycles ongoing.
  - ◆ T0 site

- ◊ CAF was down this morning for some minutes. Problem solved on the spot by alic experts [ Miguel - do you have an idea on how data rates will evolve in next few days? A: no ]
- ◊ LSF intervention - from 7th December on would be fine from ALICE side
- ◆ T1 sites
  - ◊ FZK: SE is still down. Any news? [ Xavier - now everything should be fine. Sorry could not fix yesterday already ]
- ◆ T2 sites
  - ◊ We observed some possible issues with PackMan in a couple of T2's. Software area was full and PackMan was not doing the clean up. Under investigation
- LHCb reports - MC campaign still going on, users jobs, reconstruction tail. Preparing reprocessing which should start mid next week
  - ◆ T0
    - ◊ Main issue reported in this ticket still ongoing GGUS:64043 (re-opened this morning). The plot (in full report) shows the number of user jobs failing on the grid in the last 24 hours (per site). GID is 99 = nobody instead of Z5. Why is mapping wrong? Do not play with mapping - observed problem interactively on lxplus where uid/gid properly set. Miguel - not a generic problem nor one that appeared with upgrade. Affects proxies of some users. Usually have role=lhcb. Some users coming in with no lhcb role and here mapping is failing. A small % of users and not bigger problem talking about yesterday with delays and timeouts.
  - ◆ T1 site issues:
    - ◊ RAL in scheduled downtime today.
    - ◊ Update on IN2P3 shared area ticket.

Sites / Services round table:

- KIT - nta
- NL-T1 - ntr
- BNL - unfortunately BDII issue continues to plague us. Over the course of last 24h observed that all US ATLAS resources no longer reported in SAM BDII. Even worse, resources from BNL dropped from CERN BDII. Info from CERN BDII used to dynamically install s/w needed as part of reprocessing campaign. Last night it was basically impossible to install a necessary module at BNL. Need an update from both OSG and experts from CERN to get a better handle. Ricardo - was about to update ticket; had a look again and saw problem got worse over passed few days. We see timeouts - 30" is not enough to get full info. Full OSG site is coming from 1 BDII in a single query. Putting this in reply - seems suspiciously like some tests we did when moving to SL5. Michael - you are referring to communication between OSG and CERN BDII. Ricardo - yes, CERN BDII gets info from OSG BDII. Michael - next steps? Ricardo - check response time when doing same queries from their side. Network or performance of BDII? Can increase timeout but this might not be a good solution. Also talk to Lawrence. Michael - all availability info in addition is false. All tests failing due to issues not to do with site. Ale - its a long standing issue from certain point of view the SAM responsible for experiments cannot do much more. Of course will follow up with GridView people to make sure that they follow up.
- RAL - only thing to add to Brian's report is that LHCb upgrade going well.
- IN2P3 - several points: LHCb / ATLAS shared area problem - there will be a written report to T1SCM as there is a national holiday tomorrow in Fr. CMS ticket mentioned in this meeting is in progress, especially for other experiments with similar problems. Also an occasion to say we are not so good in updating tickets but doesn't mean that people don't work here! Pre-announcement of downtime probably all day on 14th December. Tape/MSS, xrootd, FTS, LFS, other Oracle based services. Also the operations portal - downtime notifications will be delayed during outage of portal. Dashboard for EGI and NGI will still be available. More details one week before - not all completely fixed. MariaDZ - can you please put updates in LHCb tickets on shared area problems? Rolf - as I

said you will have a report but of course the tickets should be updated. Joel - one ticket has been updated just before meeting and the other put on hold as we consider more or less identical.

- ASGC - ntr
- NDGF - have dCache upgrade going on - should be about done. Found bug in dCache affecting jobs running on ARC. Only affecting a single subsite of NDGF but looking into it. Onno - which version?  
A: not sure - actually a feature missing. ARC, before getting file from cache checks if file available by getting first bytes - byte range ignored by dCache
- OSG - we got the FNAL ticketing problem fixed. So now tickets will route ok there.
- CERN Net: failover of modules in LHC OPN router at 07:30. GGUS:64126 [↗](#).
- CERN Dashboard - some problems for LHCb ATLAS and ALICE - hope all back ok tomorrow.

AOB:

- KIT - planning next major downtime - is there a schedule for next year? Will append to tomorrow's meeting.

## Thursday

Attendance: local(Eddie, Lola, Oliver, Flavia, Ricardo, Jamie, Maria, Maarten, Huang, Simone, Alessandro, Jacek, Roberto, Dirk, Gavin, Miguel, Massimo, MariaDZ, Alexei);remote(Jon, Michael Jeremy, Foued, Ronald, Joel, John, Gang, Dan Fraser).

Experiments round table:

- ATLAS reports -
  - ◆ beam setup issue last evening. restart physics fill by noon with 69 bunches beam.
  - ◆ ATLAS turns back to physics run by noon. Trigger rate increased to 300 Hz (~450 MB/s to Castor)
  - ◆ T0:
    - ◇ inconsistent throughput between GridView and ATLAS DDM dashboard (GGUS:64164 [↗](#)).
  - ◆ T1s:
    - ◇ IN2P3: disk space issue + srm highly loaded by SRMPUT (GGUS:64164 [↗](#), GGUS:64151 [↗](#)). ATLAS stopped T0 export and data consolidation to LYON and started deleting data. It's now has 100 TB space. More actions to be taken after a discussion between ADC experts and LYON people on recent SE issues at LYON.
    - ◇ TRIUMF: ATLASDATADISK is almost filled up by primary datasets (not expected by the CA-T1 share). T0 export stopped and investigation within ATLAS.
    - ◇ short SE glitch at BNL, INFN-T1 and NIKHEF last night
- CMS reports -
  - ◆ Experiment activity
    - ◇ HI data taking: no stable beam and data over night
  - ◆ CERN and Tier0
    - ◇ CASTOR load high - if this is going to happen please let us (IT-DSS) in advance.
    - ◇ noticed that castor pool cmsprodlogs fills up, deleted 3 TB, lemon does not show that the space is freed up, something wrong?  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=64175](https://gus.fzk.de/ws/ticket_info.php?ticket=64175) [↗](#)
  - ◆ Tier1 issues and plans
    - ◇ re-processing pp data + skimming at all Tier-1s
    - ◇ PileUp re-digi/re-reco requests starting

- ◇ T1\_DE\_KIT: implementation of new HI production roles:  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=64069](https://gus.fzk.de/ws/ticket_info.php?ticket=64069) -> in progress
- ◇ T1\_FR\_CCIN2P3: transfer problems to MIT, solved?  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=63826](https://gus.fzk.de/ws/ticket_info.php?ticket=63826) -> last reply to ticket from 8th Nov: "Concerning transfers exporting from IN2P3 to other sites ( like MIT who opened this ticket), We still see the same errors : AsyncWait, Pinning failed", contacted local site admins, they hope the changes they did for Atlas will solve the problems CMS is seeing (file access problems, staging issues, etc.), today is a holiday, so no update
- ◆ Tier2 Issues
  - ◇ T2\_RU\_RRC\_KI : CMS SAM-CE instabilities since several days  
(<https://savannah.cern.ch/support/index.php?117543>), site admins are following up and replied yesterday, see [https://gus.fzk.de/ws/ticket\\_info.php?ticket=63820](https://gus.fzk.de/ws/ticket_info.php?ticket=63820)
  - ◇ T2\_FI\_HIP: Wrong gstat and gridmap information for the Finnish CMS Tier-2:  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=63956](https://gus.fzk.de/ws/ticket_info.php?ticket=63956), last update Nov. 8th
    - more explanation:
      - T2 FIN resources appear in BDII as part of NDGF-T1 because we need to translate the ARC information system to Glueschema and into BDII format and this is done by NDGF
      - Site has an independent CMS dCache setup which is not part of the distributed NDGF T1 dCache setup so that cannot be published by NDGF, but is instead published by a CSC BDII server.
      - CMS applications and WLCG accounting works fine with this setup, but the WCLG monitoring is reporting non CMS resources for the T2 Fin site which obviously is wrong.
      - The ticket is about having WLCG report the right resources for us. The short term solution is to for WLCG to change configuration settings to get this corrected.
  - ◇ T2\_FR\_IPHC: several SAM tests failing:  
<https://savannah.cern.ch/support/index.php?117769>
- ALICE reports - GENERAL INFORMATION: Pass1@T0 and data replication to T1s are performing quite well. Analysis trains and three MC production cycles ongoing.
  - ◆ T0 site
    - ◇ Job submission was stopped this morning for less than 2 hours due to a problem with PackMan. Services needed to be restarted this morning in voalice13 and voalice12
  - ◆ T1 sites
    - ◇ FZK: issue Solved. Thanks
  - ◆ T2 sites
    - ◇ Usual Operation issues
- LHCb reports - MC campaign still going on, users jobs, reconstruction tail.
  - ◆ T0
    - ◇ Requested a SIR for the identified problem of xrootd re-director affecting all our user jobs since three days after the intervention on Monday. Mainly the PM must address the reason why in the risk assessment of this intervention (supposed to be transparent) it has been overlooked this potential problem. [ Miguel - need to find out where this issue is coming from. We had no change in s/w which is doing mapping based on roles. Problem with NULL roles not an incident of CASTOR upgrade. Something should be working but this problem is unrelated to upgrade. Joel - probably several issues in parallel with CASTOR. Our nightly build machinery is also affected by some problem linked to CASTOR. Probably several issues - perhaps the one in the alarm is not really linked to the intervention but sure that there is a side effect for some other activity. Have to make a clear report of all the problems on our side to

clarify. Massimo - ticket closed as getting more and more overloaded with different issues. Different issues should be followed one by one.

◇ Users are still experiencing issues accessing data (GGUS:64166 [?](#)).

◆ T1 site issues:

◇ RAL : problem accessing file with rootd (GGUS:64163 [?](#)) (authentication problem that is claimed to be fixed - checking)

Sites / Services round table:

- FNAL - ntr
- BNL - ntr for Tier1. BDII issue we discussed yesterday. We would like to update info. Extended timeout was implemented - maybe people can CERN can confirm or deny but in any case situation has slightly improved. Our resources do not drop out of BDII as often as before but still happens. May have been alleviated a bit but still there. Rob - several things on-going. Researching several potential issues. Really did see decrease in backup load on BDII's (30% load decrease) when info pared down. Been in discussions with Ricardo and others at CERN who has tested longer timeout but not implemented. Ricardo - tested on one BDII increasing timeout. Will probably do as workaround on all BDII's but need to get to root problem. Time to query remote BDII was really long. Rob - Burt has found out some latency on internet2. Plan to go and talk immediately after meeting to see if there is a network problem causing extra latency. Other thing is to validate new BDII which we want to bring into round robin (3rd BDII). Dan Fraser - one request is it is possible for some of operations team on CERN BDII to take a look at log files and may be turn on debug and see if timeouts are really issue. Ricardo - did look at logfiles and saw timeouts. OSG tickets show that timeouts increased over past few days. Dan - is that a long term fix? All: no! Michael - BNL has temporarily reduced info but cannot live without this. Temporary measure probably ok but have to have this info (s/w releases) in info system. Rob - don't know if BNL specific but ran into 5MB limit on CERN BDII's earlier this year and had to reduce that. Is there any way that BNL data can be pared down as a better set - some 3000 lines of FTS channels, some 600 of s/w release info etc. Michael - amount has been reduced to about 2.5MB. Continue to look into feasibility of reducing data. Rob - two long term fixes: network issue? more BDII's? > 1/3 of traffic is data for BNL. Rob - once we have a BDII V5 in mix can easily increase capacity.
- KIT - ntr
- NL-T1 - ntr
- RAL - ntr
- ASGC - ntr
  
- GridPP - ntr
- OSG - nta(!)
  
- CERN FTS As ATLAS previously requested the CERN-TW\_FTT and TW\_FTT-CERN channels have been added to the T0 export service.
  
- CERN LSF - deployed a new version today which should address stagein issues & maradona errors.

AOB:

## Friday

Attendance: local(Eddie, Roberto, Jacek, Massimo, Flavia, Maarten, Jamie, Maria, Huang, Ignacio, Ricardo, Oliver, Andrea A, Harry, Alessandro);remote(Rolf, Riccardo, Onno, John, Rob, Michael,Jon, Gonzalo, Tore, Xavier, Gang).

Experiments round table:

- ATLAS reports -
  - ◆ data taking for about 7 hours last night before beam tripped this morning.
  - ◆ T0:
    - ◇ ATLAS DDM central catalogue service stopped working shortly between 11:00 - 11:30 CERN time this morning. Some modification work at computer center at CERN.
    - ◇ preferable time for LFC master node hw upgrade: not earlier than 10 Dec.
  - ◆ T1s: n.t.r. [ SRM tests in IN2P3 still failing due to disk space problems reported yesterday ]
- CMS reports -
  - ◆ Experiment activity
    - ◇ HI data taking: some stable beam and data over night
  - ◆ CERN and Tier0
    - ◇ noticed that castor pool cmsprodlogs fills up, deleted 3 TB, lemon does not show that the space is freed up, something wrong?  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=64175](https://gus.fzk.de/ws/ticket_info.php?ticket=64175) -> solved and closed, advice to use stager\_rm instead of nsrm
    - ◇ talked with Castor about high load in to T0Export pool, right now rate high but ok, need to inform Castor 1/2 hour before more load will be put into pool [ Ale - what will CASTOR do if >7GB/s happens? ]
      - Computing shifters are told to call the CRC anytime when rate is going near currently observed rates
      - working out communication lines offline, not clarified yet (cell phone call, GGUS ticket of any kind?)
  - ◆ Tier1 issues and plans
    - ◇ re-processing pp data + skimming at all Tier-1s
    - ◇ PileUp re-digi/re-reco
    - ◇ T1\_DE\_KIT: implementation of new HI production roles:  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=64069](https://gus.fzk.de/ws/ticket_info.php?ticket=64069) -> in progress, last update 11/12 early morning
    - ◇ T1\_FR\_CCIN2P3: transfer problems to MIT, solved?  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=63826](https://gus.fzk.de/ws/ticket_info.php?ticket=63826) -> last reply to ticket from 8th Nov: "Concerning transfers exporting from IN2P3 to other sites ( like MIT who opened this ticket), We still see the same errors : AsyncWait, Pinning failed", contacted local site admins, they hope the changes they did for Atlas will solve the problems CMS is seeing (file access problems, staging issues, etc.)
  - ◆ Tier2 Issues
    - ◇ T2\_RU\_RRC\_KI : CMS SAM-CE instabilities since several days  
<https://savannah.cern.ch/support/index.php?117543>, site admins are following up and replied yesterday, see [https://gus.fzk.de/ws/ticket\\_info.php?ticket=63820](https://gus.fzk.de/ws/ticket_info.php?ticket=63820)
    - ◇ T2\_FI\_HIP: Wrong gstat and gridmap information for the Finnish CMS Tier-2:  
[https://gus.fzk.de/ws/ticket\\_info.php?ticket=63956](https://gus.fzk.de/ws/ticket_info.php?ticket=63956), last update Nov. 8th, further discussed in T1 coordination meeting, Flavia following up as well
      - more explanation:
        - T2 FIN resources appear in BDII as part of NDGF-T1 because we need to translate the ARC information system to Glueschema and into BDII format and this is done by NDGF
        - Site has an independent CMS dCache setup which is not part of the distributed NDGF T1 dCache setup so that cannot be published by NDGF, but is instead published by a CSC BDII server.
        - CMS applications and WLCG accounting works fine with this setup, but the WCLG monitoring is reporting non CMS resources for the T2 Fin site which obviously is wrong.
        - The ticket is about having WLCG report the right resources for us.

The short term solution is to for WLCG to change configuration settings to get this corrected.

- ◇ T2\_FR\_IPHC: several SAM tests failing:  
<https://savannah.cern.ch/support/index.php?117769>, no update since opened on 11/11
  - ◇ T2\_CH\_CSCS: PhEDEX FileDownload agent down:  
<http://savannah.cern.ch/support/?117805>
  - ◇ T2\_RU\_JINR: SAM CE test fails analysis check, file not accessible:  
<https://savannah.cern.ch/support/index.php?117804>
  - ◇ T2\_US\_Nebraska: unscheduled downtime due to power cut
- ALICE reports - GENERAL INFORMATION: Pass0 and pass1 reco and data replication to T1s are performing quite well. Depending on the LHC luminosity we expect to reach the standard data taking rate of 1.5GB/sec (2.5GB/sec max) for next week.
    - ◆ T0 site
      - ◇ ntr
    - ◆ T1 sites o ntr
    - ◆ T2 sites
      - ◇ Usual Operation issues
  - LHCb reports - Massive clean up campaign launched across all T1 sites for old production and user spaces.
    - ◆ T0
      - ◇ Debugging the issue with accessing data at CERN. If the problem is still there on Sunday evening, we ask to roll back on monday morning, the intervention of last Monday, even if, in principle, there is no relation with the problem. At least we will have a clear proof of it. [ Massimo - "CASTOR" problems boils down to issue seen before upgrade. Certificate seen by server is "corrupted". Proxies on WNs go through a complicated path - has to be understood. Have a kind of work-around - detecting "corrupted" proxies and falling back to normal user. Not guaranteed to work. Ignacio - ignore voms map file. ] **Before any such action a meeting between IT-DSS and LHCb is strongly recommended, e.g. on Monday morning.**
    - ◆ T1 site issues:
      - ◇ GRIDKA: observed ~80% of the transfers to the LHCb\_MC-M-DST space token (connection time out error). Under investigation, GGUS ticket eventually to be submitted by the shifter.

Sites / Services round table:

- IN2P3 - still working on ATLAS/LHCb/CMS problem. Won't add anything to report from yesterday (written). Alessandro - thanks. Lyon and whole FR cloud will disable all activity R/W Lyon (except reprocessing).
- BNL - ntr for T1. Leave to Rob to report on BDII progress. The network engineers still working on CNAF to BNL network problem. Upgrading MCC in Amsterdam - also movement there.
- OSG - as far as BDII BNL issues we are following on several directions. Timeout increased to 120s at CERN (Ricardo). This should ease problem for now. Been working with internet2 engineers to locate possible problem on network. Have seen network saturation at some times. 3rd item: have a new version of BDII that is ready to put into round robin which would alleviate some of load and hope to have permanent solution to network latency and/or capacity and allow CERN to reduce timeout and BNL to publish full info. Flavia - as reported to you Lawrence asked site admin in Australia to see if BNL info appears in Australia. It does not. Is network path the same? Rob - good to have traceroute to see if path has internet2 component. Handover between internet2 and local campus? Haven't eliminated any possibilities. Tim Dyce in Australia to be contacted for traceroute.
- CNAF - ntr

## WLCGDailyMeetingsWeek101108 < LCG < TWiki

- NL-T1 - reminder: Tuesday we will have downtime at SARA - 08:00 - 18:00 - announced in GOCDB because of router maintenance also affects LHC OPN. GGUS ticket submitted: [GGUS:64142](#)
- RAL - yesterday we replaced LHCb SRM server - went from 2 to 3 machines. Next week outage for CASTOR CMS for upgrade to 2.1.9 (Tue-Thu)
- FNAL - ntr
- PIC - yesterday at around 14:00 UTC had a 30' glitch which affected all T1 servers - problem on CRL of Spanish CA. All symptoms pointed to CRL expiry - happens once a year on the day that it is renewed. Automatically solved when proper CRL is in place. A lot of ATLAS transfers failed. Got GGUS ticket but solved quickly.
- NDGF - ntr
- KIT - one of f/s setup this week for increase of ALICE storage area continuously running into kernel panics
- ASGC - ntr
  
- CERN - ntr

AOB:

-- JamieShiers - 04-Nov-2010

---

This topic: [LCG > WLCGDailyMeetingsWeek101108](#)

Topic revision: r13 - 2010-11-12 - JamieShiers



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.  
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)