

# Table of Contents

<b>Week of 091123.....</b>	<b>1</b>
WLCG Service Incidents, Interventions and Availability.....	1
GGUS section.....	1
Daily WLCG Operations Call details.....	1
Monday:.....	1
Tuesday:.....	3
Wednesday.....	5
Thursday.....	6
Friday.....	7

# Week of 091123

## WLCG Service Incidents, Interventions and Availability

VO Summaries of Site Availability				SIRs & Broadcasts	
<a href="#">ALICE</a>	<a href="#">ATLAS</a>	<a href="#">CMS</a>	<a href="#">LHCb</a>	WLCG Service Incident Reports	Broadcast archive

### GGUS section

- GGUS Escalation reports every Monday (used for WLCG Service Report to MB)
  - ◆ Full search: [https://gus.fzk.de/ws/ticket\\_search.php](https://gus.fzk.de/ws/ticket_search.php) and select "VO" - this will give all tickets, including team & alarm
- LHCOPN Tickets in GGUS: see [https://gus.fzk.de/pages/all\\_lhcopn.php](https://gus.fzk.de/pages/all_lhcopn.php) and change your selection criteria. Future actions are also listed.
  - ◆ If network group participation is necessary, please invite them in time.
- OSG items selected from the GGUS escalation reports. This week, report from Guenter Grein:

All tickets are synchronized besides [https://gus.fzk.de/ws/ticket\\_info.php?ticket=53365&from=search](https://gus.fzk.de/ws/ticket_info.php?ticket=53365&from=search) which seems not to make it into OSG system until now, although the mail was sent from GGUS to OSG on 2009-11-18 10:09:09 UTC.

- Procedure to become a LHC Experiment VO TEAM member
- Other recent GGUS FAQs

### Daily WLCG Operations Call details

To join the call, at 15.00 CE(S)T Monday to Friday inclusive (in CERN 513 R-068) do one of the following:

1. Dial +41227676000 (Main) and enter access code 0119168, or
2. To have the system call you, [click here](#)
3. The scod rota for the next few weeks is at ScodRota

General Information			
CERN IT status board	M/W PPSCoordinationWorkLog	WLCG Baseline Versions	Weekly joint operations meeting minutes

Additional Material:

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

### Monday:

Attendance: local(Simone, Alessandro, Miguel, Jan, Harry, JPB, Olof, Patricia, Roberto, Guiseppe, Maria, Maria, Dirk); remote(Daniele/CMS, Gonzalo/PIC, Ono/NL-T1, Alexei/ATLAS, Rolf/IN2P3, Angela/KIT, Jeremy/GridPP, Jason/ASGC, Gareth/RAL, Michael/BNL, Kyle/OSG Paolo/CNAF).

Experiments round table:

- ATLAS - (Alexei) ATLAS is taking LHC data and exports to T1/2/3 (60 centers) smoothly. Big thanks to Miguel for his help during a SRM problem Fri afternoon (details below). During this problem also GGUS alarm system did not work. Dirk: a service incident report is requested from the

GGUS responsible. ATLAS sees some T1-T1 transfer problems: FZK has (had) dCache problem. Data transfers from FZK to BNL and from FZK to T2s were affected. More info will be available after FZK meeting later today. There also was competition between data and user datasets transfer from BNL to LYON (in addition LYON had problem with one of the dCache servers). ATLAS also experienced larger spikes in DB access on Sun. Luca Canali noticed spikes at 9:00 am and 30 mins later ATLAS notices slow DB access. Investigation between DB responsible in ATLAS and IT initiated. Simone: SRM issue on Fri 5:30-6:00pm with increased in transfer error rate to all T1s. Symptom looked similar to previous problem and are not related to SRM load. FTS started showing source errors, as GGUS was not operational experiment called castor operations directly and the problem was fixed in 10 mins by a SRM restart. No impact in data taking and the problem did not reappear on weekend. Still ATLAS has during the last period seen several SRM related problems and request a plan to find the root cause. Miguel: also Jan and Giuseppe were looking at the problem when it occurred. In the meantime a service incident report has been produced with the state of the analysis and an action plan to find the root cause and several other suggested improvements to the service. ATLAS suggests to send another test alarm ticket to confirm that the system is working for them. Depending on the problem analysis from the GGUS team, this should probably also be tested for the other experiments.

- CMS reports - (Daniele) On 20 nov 19:00 - beam at P5, splash events, possibly collisions but high halo rate. Status of all tickets up to date in the CMS twiki. CMS is was affected by a condor bug (affected MC) - received working patch, more detailed report at 5pm CMS meeting. CMS lost 100 TB of custodial data at IN2P3 due to a communication problem between CMS and site staff. Operation procedures being reviewed now.
- ALICE - (Patrica) stressing transfer infrastructure with new alien version, look very promising so far. Some T2 related issues over the weekend being followed up directly with the sites. Had to reopened ticket concerning CREAM CEs at CERN after problems - being worked on.
- LHCb reports - (Roberto) collisions/splash events recorded also in LHCb - first files fully reconstructed and shipped. Problem at NIKHEF: WN jobs failing due to stuck nsd daemon (Root fails as it can not figure out \$HOME). Site has put an nsd watch dog put in place. SARA: users with double VO cert (ATLAS and LHCb) failing as since the dCache upgrade to the golden release - dcap seems to be back at static VO mapping and certs were mapped to ATLAS. Now statically mapping to LHCb as temporary workaround. Problem has been reported to dCache developers who work on a fix. PIC: run out of space for shared area for Brunel installation. PIC added more space. CNAF: On Fri STORM went down but was restarted quickly. All issues covered by GGUS tickets.

Sites / Services round table:

- Gonzalo/PIC: ntr
- Ono/NL-T1: NIKHEF problem with LHCb - workaround in place. nsd is checked every 1 min. Is rate sufficient? Roberto: when was this put in place. Users still complained on Sat. To be followed up with between LHCb and NIKHEF. NIKHEF: new computing nodes had to be put offline as communication problems with storage started appearing during this weekend - under investigation. SARA: SRM problems reported on Fri were not caused by iptable config: root cause was SRM dcache crash (confirmed as 1.9.5-8 bug due to use of non-thread safe library). Now fixed in 1.9.5.-9 which has been installed and works fine. Site is working on new storage h/w setup which should become available tomorrow or Wed.
- Angela/KIT: Some discussions on ATLAS team ticket: local contact considered this rather an ALARM ticket. Simone: no, thats fine - shift decides on severity.
- Garaeth/RAL: Rescheduled at risk time for bypass test for tomorrow morning (being added to

GOCDB).

- Jason/ASGC: new resources delivered to the site.
- Michael/BNL: one storage server down for 45 min - restored. Simone: concerning transfer backlog from FZK to BNL - observed long latency for turl at FZK. Can you confirm that BNL SRM is 2.2 in legacy mode? BNL will get back to ATLAS.
- Paolo/CNAF: ntr
- Rolf/IN2P3: "at risk" scheduled for Wed to split EGEE VOs from LHC VOs. Should be transparent for LHC VO s Alessandro: how long is the at risk? ->All day.
- MariaG/CERN: high load on ATLAS DB on weekend - first DB node 3 (panda -9:15 evening) then node 5 (dq2 - 9:20). will follow up the cause of the large spikes (factor 5-6 of maximum load) with ATLAS DB experts. Other nodes of the cluster stayed unaffected (eg conditions). Also continuing plan to moving other application off the atlas cluster and to rebalance core apps (eg isolate panda and dq2 on dedicated DB nodes).

AOB:

- Simone: DDM node of central catalog lost DB connection - due the cx\_ora problem fixed in SL5. Question: Would be handy to take out node from DNS load balancing as workaround/ to simplify problem analysis. Can we have a procedure to reconfigure DNS alias for a small set of experiment experts?
- MariaD: Note to Kyle: ticket 53365 GGUS is not synchronized with OSG - please have a look and add info to ticket.

Further info in the ticket itself.

- **ATLAS ALARM ticket by Alexei:** In reply to the post-mortem request at the wLCG meeting today, the solution in [https://gus.fzk.de/ws/ticket\\_info.php?ticket=53520](https://gus.fzk.de/ws/ticket_info.php?ticket=53520) by Guenter Grein contains the full explanation. We hope this suffices.
- **GGUS fail-safe system:** MariaD had a phone conversation after today's wLCG meeting with the GGUS team leader Torsten Antoni. Updates in <https://savannah.cern.ch/support/?101122> The issue will be discussed at the next USAG meeting on December 10th <http://indico.cern.ch/conferenceDisplay.py?confId=73657> and at the MB tomorrow.
- A service incident report for the ATLAS SRM problem on Fri has been produced by the CERN castor team. Please refer to <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGServiceIncidents>

## Tuesday:

Attendance: local(Olof, Simone, Jamie, Nick, Jean-Phillipe, Alessandro, Roberto, MariaG, Harry, Andrew Gavin,Daniele, Evan, Patricia, Steve, Julia, Gang, Wei, Dirk);remote(Alexei/ATLAS, Michael/BNL, Gonzalo/PIC, Rolf/IN2P3, Xavier + Angela/KIT, Gareth/RAL, Ronal/NL-T1, Angela/KIT, Jeremy/GridPP, Brian).

Experiments round table:

- ATLAS - (Alexei) running smoothly, export to T1,2 - no major issues since yesterday. RAL CASTOR DB issues required memory increase. Simone: discussion with castor@cern, went through

Tuesday:

post mortem proposals - now running with tuned setup suggested by castor team. Shifters will monitor and give feedback on improvement. Data distribution: reverted to previous distribution workflow after concerns from T1 were raised about T1-T1 transfer load.

- CMS reports - (Daniele) CMS is exporting to T1 and monitoring replication progress to T2. Validating phedex priorities for data (vs MC). two issues closed: CERN - FNAL transfers are back to low error rate since the changes on the CASTOR and dCache side. Also followup on condor bug (discussed at yesterday's CMS ops meeting) concluded - the applied patch resolved the issue (back to few percent errors). No major problems at T1. Recovery from IN2P3 data deletion done - repopulating the site and expect it back in full production soon. T2: some issues with CMS SW installation - details on the CMS twiki.
- ALICE - (Patricia) 5:30 first p+p run went successfully through full processing chain: castor, reconstruction, export to five sites. Success also for grid infrastructure. In addition MC production (currently 5k jobs) ramping up. No outstanding issues at tier sites.
- LHCb reports - (Roberto) Also LHCb had 56min after the first collisions data fully reconstructed (registered in their bookkeeping, shipped to T1, redistributed output). Big achievement! No MC production right now. Significant analysis activity visible via the dashboard - triggered by first data. FIO confirmed they are ready to setup the requested LHCb\_HISTO space token and at the same time to allocate the 100TB of disk space now ready to be installed. Up to LHCb to decide how to spread this extra space across various service classes. VO card for shared area updated to 150GB to cope with more platforms and retain a safety margin at all sites. ALARM ticket for DB streaming from CERN. All T1 DBs but CNAF affected. Physics DB team opened Oracle service request and improved streams monitoring. Reestablished streams for Conditions, but leaving LFC replicas (currently unused) for problem analysis by Oracle. NIKHEF problem (ROOT unable to obtain \$HOME location) reported still on Sat was due to new WN which moved into production.

Sites / Services round table:

- Gonzalo/PIC : ntr
- Xavier/KIT: dcache disk only pools became unavailable (no read access). Workaround by marking the pool read-only only in the pool manager but not the pool itself. Another pool reconfig took place to increase the max number of inodes.
- Gareth/RAL: planned at risk for UPS tests postponed due to LHC activity. CASTOR ATLAS: performance problems on SRM DB yesterday evening and this morning. Reduced FTS channels to reduce SRM activity and error rate. Now resolved by moving Oracle processes to a node with more memory. Q: can we monitor the backlog from ATLAS? Simone: will send out a collection of links to monitoring pages and some comments on how to interpret information (eg open data sets). Gareth pointed out that current list of FTM endpoints is incorrect. Gavin: will fix the list.
- Rolf/IN2P3 - ntr
- Ron/NL-T1: first new storage node available, added to space tokens for ATLAS. More to come soon also for other VOs. New WNs back in production: recent network problem was found to be due to broadcom network driver feature (transparent packet aggregation) - feature now disabled.
- Jeremy: ntr
- Michael/BNL: ntr
- Gang/ASGC: seeing low efficiency of tape media use - what are typical values from other sites. Will ask the question on the castor external list. Additional statement by email from Jason: castor tape i/o problem with new installed LTO4 tapes. will extend the discussion with support team later on. seems more related to tape packages for the LTO4 support, and need to validate the check sum from tape dump. Online offline priority still being discussed with experiments.
- Gavin/CERN: would like to rename FTS 2.2 endpoint (currently FTSPPS) this week
- MariaG/CERN: Atlas offline cluster: small instabilities due to adding a fifth node . More details on

LHCb for streams problem: apply of streams did not process updates (4 sites affected since their last intervention). Oracle process monitoring did not show any problem. After the DB team was notified conditions were resynchronized within a few hours. GGUS alarm tickets for databases: will need procedures for operators. DB team can also be contacted via phydb.support and on-call phone. Harry: has DB piquet already started? not yet. Alessandro: 24h coverage for DBs? Currently best effort - but soon piquet.

AOB: report by email from Steve:

### **CERN LHC VOMS Service was unavailable for 20 minutes**

Monday 16:02 -> 16:21 UTC there was complete unavailability of the voms service. Cause was a misconfiguration of the gd lcg-fw service. Service restored after automatic notification to service manager.

## **Wednesday**

Attendance: local(Olof, Daniele, Jean-Phillipe, Simone, Jason, MariaD, MariaG,Patricia, Andrew, Roberto, Julia, Gang, Carlos, Lola, Alessandro, Antonio, Dirk);remote(Xavier/KIT,Michael/BNL, John/RAL,Tore/NDGF, Rolf/IN2P3 , Ron/NL-T1, Paolo/CNAF, Alexei/ATLAS ).

Experiments round table:

- ATLAS - (Alexei) Reprocessing of beam data. Smooth operation, with some transfer problems: BNL->CERN and BNL->IN2P3 (either FTS channel config or network)- all relevant experts involved
- CMS reports - (Daniele) closing MC prod tickets, some transfer errors FNAL -> IN2P3 (timeout with larger files needs retuning), data distribution followed up down to t3, smooth running, still some margins for optimisation. IN2P3 are back in production and now get minimum bias events. T2: closing Bari - new instabilities for SAM test concerning CMS SW installation at some sites.
- ALICE - (Patricia) CE201 not reported in BDII due to an issue with the information provider - can't be used yet. KIT: discrepancy between DB and BDII information. Solution: manual cleanup of CREAM db. Alice will check if the sites this has performed this cleanup. Alice checks for backup for all VO specific services. Introduced voalice11 as backup of cream CE which replace production node incase of problems (manual intervention). New issue which could affect several sites: registration of user proxy in vobox due to discrepancy between vobox and alien library versions - alien code side tested and is working fine.
- LHCb reports - (Roberto) LHCb mainly running analysis - 6k jobs a day (not much). Special MC production for current beam conditions in preparation. No major new issues. NIKEF issue with nscd/ROOT: no more user complains. Lyon: file access issue now understood: dcap was not running in active mode on some machines. Now fixed.

Sites / Services round table:

- Xavier/KIT: ntr
- Micheal/BNL: ntr
- John/RAL: ntr
- Tore/NDGF: ntr
- Rolf/IN2P3: at risk for LHC experiment services, but intervention is assumed to be transparent.
- Ron/NL-T1: ntr
- Paolo/CNAF: ntr
- Gang/ASGC: ntr

- MariaG/CERN: rolling DB intervention at SARA - some DB session were relocated. Load now back to normal.
- Antonio/CERN: s/w versions: not many news : Release report: deployment status wiki page Simone: FTS 2.2.2 timescale? Antonio: not clear yet.

AOB:

- MariaG: A distributed database workshop will take place at CERN in the next two days: <http://indico.cern.ch/conferenceOtherViews.py?view=standard&confId=70892>

## Thursday

Attendance: local( Olof, Daniele, Harry, Jean-Philippe, Evan, Andrew, Jan, Roberto, Gang, Simone, MariaD, Dirk);remote(Tore/NDGF, Angela/KIT, Gareth/RAL, Brian/GridPP, Elisabetha/CNAF, Ronald/NL-T1).

Experiments round table:

- ATLAS - (Simone) transfers in BNL-SARA-LYON triangle - two issues: traffic BNL->SARA and BNL->IN2P3 was looked by experts from three sites. BNL-SARA was improved by increasing FTS transfer slots from 10 to 30 - now reaching 160MB/s. Second issue (BNL->IN2P3) not yet fixed: 20 transfer active but less than 1 MB/s, FTS logs look fine but 20% of transfers are killed in transfer stage. Points to gridftp problem but iperf tests not yet conclusive. ATLAS opened ticket against IN2P3. ASGC: low rate on T1->T2 transfers. Many active transfers jobs in FTS, but long time to process ( 200 transfers took 5)h. ticket open to ASGC, now rate improved 50 MB/s. As ASGC has 50% of data and the backlog is increasing this could become a problem. ASGC: believe to have found and fixed the problem yesterday 23:00. ATLAS will check.
- CMS reports - (Daniele) not much to highlight, details for T2 followup on CMS twiki. More and more T2 involved. Tier1: now 6 (of 7) sites being populated, priority given to custodial sites. Data subscriptions vary and not all data taken from CERN. Some examples mentioned by Daniele: CERN to FNAL peaks at 500MB/s - allows redistribution to US sites. Some problems due to files >10GB (FTS timeouts). Proxy related problems at some T1. Also good performance on T2 channels: eg FNAL to San Diego 300MB/s, FNAL to Florida at 100 MB/s. Transfer quality generally good.
- ALICE - Info provider of CE201 is dead since Tue: as also CE202 is currently being reinstalled, ALICE can not use CREAM CE at CERN. FIO looking into this problem.
- LHCb reports - (Roberto) No special issues - low activity - mainly analysis (3k jobs). CNAF has been suggesting a castor SRM upgrade - LHCb discourages upgrades before xmas pause if not related to significant service problems. LHCb has marked CREAM CE as critical on their dashboard.

Sites / Services round table:

- Tore/NDGF: ntr
- Angela/KIT: yesterday evening pbs problem (wall time limit reached) - recovered by itself. Also many FTS transfers in status hold, monitoring did not catch this problem. Now monitoring and root cause fixed. CMS and LHCb have affected.
- Gareth/RAL: ntr
- Elisabetha/CNAF: ntr
- Ronald/NL-T1: ntr
- Gang/ASGC: ntr
  
- Jan/CERN: propose upgrade of SRM: Mon for client, Wed: server (LHCb, CMS and ATLAS) - experiment feedback welcome - email proposal will be sent

- MariaD/CERN : from GGUS developers: GGUS will show on all tickets the GOCDB site name and (in parenthesis) the WLCG name. Help for searches will be provided (even though GGUS DB will contain only GOCDB name). People familiar with the WLCG name only submitting/searching GGUS tickets to T1s (e.g. alarm tickets) should be informed. Relevant savannah item <https://savannah.cern.ch/support/?110614> Second suggestion from GGUS developers: would like to do (more) regular alarm tests triggered directly by GGUS team and would need privileges to submit test tickets from the VOs: ATLAS and CMS agree. T1 sites will be informed about this proposed (monthly) test at the next MB to get their feedback. Dirk will bring it to the MB next Tuesday. Relevant savannah item: <https://savannah.cern.ch/support/?111475> Daniele: CMS finished defining the list of people who can submit TEAM tickets in VOMS and tested it successfully.

AOB:

## Friday

Attendance: local(Lola, Jan, Nick, Jean-Phillipe, Roberto, Simone, Julia, Dirk);remote(Jeremy/GridPP, Angela/KIT, Rolf/IN2P3, John/RAL, Tore/NDGF, Gang/ASGC,Onno/NL-T1).

Experiments round table:

- ATLAS - (Simone) slow transfer ASGC T1->T2 - fixed now. Reason was small clock skew of 10/s which affected FTS. FZK: 0.5-1h failures yesterday evening due to error locality unknown. BNL->IN2P3 still transfers at low rate (1MB/s). IN2P3 in touch with BNL and checking if problem is bi-directional. Needs fixed soon otherwise backlog will build up. Most critical: FTS 2.2 is not delivering files from CERN. Largest backlog to BNL (50 jobs - 250 files are ready but not activated in FTS), but also CERN to RAL (transfer waiting since many days). If this issue can not be fixed quickly then ATLAS would be prepared roll back to FTS 2.1

(need additional discussion). Jan: FIO is looking into this - may be related to recent machine move. Gavin: other FTS agents seem ok, one agent is polling SRM repeatedly and created already large log file. Developers are looking at core files. Simone: need feedback/decision soon. FZK on locality issue: only see problem on source (UNI Wuppertal) - UNI W should check. Yesterday evening problem has been tracked down to a firewall problem.

- CMS reports - Daniele had to apologize but the CMS twiki has been updated.
- ALICE - (Lola) CE201 was back in information system and ALICE production, but new failures this morning and out of production again: no CREAM CE available at CERN to ALICE.
- LHCb reports - (Roberto) no activity right now. Acknowledge that the recent SARA upgrade now fulfills the LHCb storage requirements.

Sites / Services round table:

- Rolf/IN2P3: ntr
- Angela/KIT: after discussion with local ATLAS representative: downtime proposed for next tue (2h) - all nodes will be upgraded to 1.9.5 or higher. Glexec is now available at KIT.
- John/RAL: ntr
- Tore/NDGF: ntr
- Onno/NLT1: ntr
- Gang/ASGC: ntr

AOB: \* Eduardo/CERN: this morning 9:00 for 1h heavy transfers to GSI (2GB/s) overloaded firewall - ATLAS will follow up. \* Simone/ATLAS: ATLAS decided to skip the proposed SRM intervention. Jan:

Friday



ALICE agreed, waiting for feedback from LHCb and CMS.

-- JamieShiers - 20-Nov-2009

---

This topic: LCG > WLCGDailyMeetingsWeek091123

Topic revision: r12 - 2009-11-27 - DirkDuellmann



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