

Table of Contents

Week of 120604.....	1
Daily WLCG Operations Call details.....	1
WLCG Service Incidents, Interventions and Availability, Change / Risk Assessments.....	1
General Information.....	1
Monday.....	1
Tuesday.....	3
Wednesday.....	5
Thursday.....	6
Friday.....	8

Week of 120604

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
ALICE	ATLAS	CMS	LHCb	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(Luc, Jamie, Simone, David, Maria, Ignacio, Jan, Maarten);remote(Kyle, Joel, Lisa, Jhen-Wei, Giovanni, Roger, Onno, Ian, Pavel, Rolf).

Experiments round table:

- ATLAS reports -
- LHC/ATLAS - physics data taking. Important and urgent MC production on-going (not new).
- WLCG services
 - ◆ Alarm ticket GGUS:82797. ggus authentication failing with CERN CA. Solved.
- T0
 - ◆ Alarm ticket GGUS:82811. Very slow LSF response time (bsub, bresume). System update to batch slowed the masternodes. Solved.
- T1
 - ◆ Alarm ticket GGUS:82791 SARA transfers failing (destination). NL cloud broker off Sat & Sun. SARA out of T0 export. NL cloud has been put online again, SARA had downtime for new h/w
 - ◆ INFN-T1 Stable. Fully back in T0 export (disk & tape).

- CMS reports -

- LHC machine / CMS detector
 - ◆ Physics ongoing.
 - CERN / central services and T0
 - ◆ Problem with LSF earlier today, which exposed an issue in the CMS system. If the workflow submitters cannot tell the number of jobs, they submit more.
 - Tier-1/2:
 - ◆ Problems with ASGC storage. (unexpected downtime)
 - ◆ Backlog at RAL cleared
 - Other
 - ◆ CRC stays Ian Fisk
-

- ALICE reports [↗](#) -
 - ◆ KIT (most important T1 for ALICE) VOBOX unusable due to NFS problem since Sat afternoon; being fixed [Pavel - is problem on ALICE side or on KIT side? A: remaining problem is on ALICE side]
-

- LHCb reports -
- Users analysis and prompt reconstruction and stripping at T1s ongoing
- MC production at Tiers2

- **New GGUS (or RT) tickets**

- T1:
 - ◆ IN2P3: 26k files unavailable due to hardware problem (GGUS:82751) the same disk server again failing [Rolf - as signalled first time cause for failure of diskserver not known and hence upgraded microcode for next crash. Manufacturer now recommends to change mother board and 2 CPUs of machine. M/C currently up and running - please decide whether you want to access files whilst we copy to another m/c or else you will blacklist them. During copy m/c might crash!
 - ◆ SARA : downtime between 9h and 12h UTC. Late announcement. Please open relevant ticket... (OPS portal)
-

Sites / Services round table:

- PIC: Yesterday June 3rd at around 18:00 CEST PIC suffered a power glitch which caused the cooling of part of the infrastructure to halt. The resulting overheat forced us to shutdown part of the WNs to avoid further problems. About 2500 WNs were powered off. The WNs have been powered on and put back in the production Computing service this morning at around 12:00 CEST. The service is now fully restored. We will prepare a SIR to analyse the technical issues of the infrastructure incident.
- FNAL - ntr
- ASGC - this morning we had some problems whilst doing CASTOR DB maintenance. Could not be solved before downtime over, now another unsched downtime.
- BNL - ntr
- CNAF - ntr
- NDGF - Friday report on tape copies; as far as we can tell no files have been lost and all tape drives up and running
- NL-T1 - SARA Downtime finished at 15:00 CEST ; had to extend by an hour because more work than foreseen. dCache namespace DB now running on faster h/w and hope that this will solve stability issues; please let us know!
- KIT - ntr

- IN2P3 - nta
- OSG - ntr
- CERN storage - looking into recent ATLAS ticket about 60+ lost files. GGUS:82826
- CERN dashboard - FTS monitoring; last Friday had confirmation for all sites that ActiveMQ patch was deployed; still not getting data from KIT - will follow up offline between developer and site

AOB: (MariaDZ) File ggus-tickets.xls with total numbers of GGUS tickets per experiment per week is up-to-date and attached to WLCGOperationsMeetings page. There were two real ALARMS last week GGUS:82791 for SARA-MATRIX and GGUS:82797 for KIT (for GGUS web pages' unavailability for users with CERN certificates). Both ALARMS were raised during the weekend of 2012/06/02 - 17 real ALARMS in total since the last MB. Slides with detailed drills are attached to this page.

Tuesday

Attendance: local(David, Eva, Gavin, Ian, Ignacio, Luca M, Maarten, Maria D, Nilo, Simone);remote(Giovanni, Gonzalo, Jhen-Wei, Lisa, Michael, Onno, Rob, Roger, Rolf, Xavier M).

Experiments round table:

- ATLAS reports -
 - ◆ T0
 - ◇ Alarm ticket GGUS:82839. Very slow LSF response time (bsub, bresume). Different cause w.r.t. the issue reported yesterday. Today the cause is massive amount of pending jobs in the system. Some power user in ATLAS has been warned and is being educated. A protection from the LSF side would be more than welcome.
 - Gavin: today's problem not yet understood, but correlation with high query load; ticket opened with Platform, maybe a bug in the latest version; response time improved after some user scripts were stopped; experiments should try reducing query load and smoothing the submission rates
 - ◇ Several files could not be exported from CASTOR to T1s yesterday afternoon. The error indicates a problem with the disk server ("no copies available"). The file is also on tape but I am not sure CASTOR tries to recall it. GGUS:82830. All files made it after 6PM.
 - ◇ Since about 13:00 we are having problems both in retrieving data from and writing data to CASTOR. Both T0ATLAS and T0MERGE are affected. ALARM has been sent. GGUS:82854
 - Luca: all experiments were affected by contention in the Name Server DB which appears to have been caused by the SRM; it went away by itself
 - Nilo: we saw a very high Name Server load coming from ongoing deletions; the load dropped when those activities stopped; to be understood what went wrong with those operations - deletions are standard operations that normally do not pose problems
 - ◆ T1
 - ◇ SARA is still trouble. After the upgrade yesterday and some follow up intervention, we still observe 25% error rate (at least). The usual ticket GGUS:82490 has been updated. We still have not clear if the issue is overload and who causes the overload. If the overload comes from ATLAS, please provide IP and DN of the aggressive client. If the overload comes from someone else, please protect ATLAS from someone else. SARA partial unavailability (today lasting since more than one month) is causing many troubles to the preparation for ICHEP which is one of the main milestones of the year.

- Onno: Gerd Behrmann of the dCache team suggested a field type mismatch in a foreign key constraint may cause extreme delays when deleting rows, because indexes would not be used. NDGF has had a similar problem a while ago. We have such a field type mismatch in the table that contains the ACL information; we are in a downtime now to correct that and to re-index the whole DB at the same time. We will consider giving ATLAS and LHCb their own instances at some point, but that would require very careful planning because 1 or 2 PB would have to be moved to dedicated pool nodes (now they are all shared).
- ◇ One RAW file still not replicated to ASGC after the CASTOR Oracle intervention. News on this?
 - Jhen-Wei: CASTOR was back this morning and successful transfers were seen; will check what happened to that file
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ Physics ongoing.
 - ◆ CERN / central services and T0
 - ◇ Problem with LSF earlier today. Slow response from LSF. CMS trying to reduce the query rate also.
 - GGUS:82845 [↗](#)
 - Ian: CMS currently submit more jobs when the response time is slow, while the code should back off instead
 - ◆ Tier-1/2:
 - ◇ Problems with ASGC reported to be solved.
- ALICE reports [↗](#) -
 - ◆ Central services partially down between 8:00 and 10:00 CEST due to HW problems with UPS and 1 machine.
 - ◆ KIT working again since 16:00 yesterday.
- LHCb reports -
 - ◆ Users analysis and prompt reconstruction and stripping at T1s ongoing
 - ◆ MC production at Tiers2
 - ◆ **New GGUS (or RT) tickets** [↗](#)
 - ◆ T1:
 - ◇ IN2P3: (GGUS:82751 [↗](#)) on going

Sites / Services round table:

- ASGC - nta
- BNL
 - ◆ since noon UTC connected to LHCONE for ATLAS computing, accepting prefixes from ~60 participating networks
- CNAF - ntr
- FNAL - ntr
- IN2P3 - ntr
- KIT -
 - ◆ Could not join phonecon, "conference not started"?
 - ◇ see AOB
 - ◆ Problems with one tape lib, which is halting frequently. Called in external support for tomorrow.
- NDGF - ntr
- NLT1 - nta

- OSG - ntr
- PIC - ntr

- dashboards - ntr
- databases - ntr
- GGUS/SNOW - ntr
- grid services - ntr
- storage
 - ◆ CASTOR problem 12:45-13:45 CEST affecting all experiments, see ATLAS report

AOB:

- SNOW ticket INC:134824 [↗](#) opened about today's Alcatel problem ("meeting has not started")

Wednesday

Attendance: local (Andrea, David, LucaC, Maarten, Simone, Jan, MariaD); remote (Gonzalo/PIC, Lisa/FNAL, Roger/NDGF, Tiju/RAL, Jhen-Wei/ASGC, Ron/NLT1, Rolf/IN2P3, Rob/OSG, Dimitri/KIT; Joel/LHCb, Ian/CMS).

Experiments round table:

- ATLAS reports -
 - ◆ T0
 - ◇ NTR
 - ◆ T1
 - ◇ SARA seems out of the tunnel. No particular issue observed this morning (or, better, after the intervention yesterday). Will keep watching.
 - ◇ MCTAPE endpoint at TRIUMF have been set to "write-only" mode, following the request of the site (there is a planned intervention).
 - ◇ Some transfers to TRIUMF are failing. The problem seems to be very ":source-selective" but there is no obvious correlation with network paths. GGUS:82872 [↗](#).
 - ◆ OTHERS
 - ◇ GGUS support contacted for a problem with TEAM tickets submitted by comp at P1. Those tickets appear with VO=none, so they are not in the list of ATLAS open GGUS. [Maarten: is this problem at P1 recent? Simone: from this morning. Maarten: instead of using the mailing list, it would have been better to open a GGUS ticket about this issue.]
- CMS reports -
 - ◆ LHC machine / CMS detector
 - ◇ Physics ongoing.
 - ◆ CERN / central services and T0
 - ◇ Clean-up ongoing from LSF issues. Seem to have lost track of some jobs with the slow response of Tuesday
 - ◆ Tier-1/2:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ Users analysis and prompt reconstruction and stripping at T1s ongoing

- ◆ MC production at Tiers2
- ◆ T0:
 - ◇ CERN: (GGUS:82874) SRM BUSY for transferring file from PIT to CASTOR * T1:
 - ◇ IN2P3: (GGUS:82751) Fixed
 - ◇ Set Inactive all the readonly instance of LFC in order to prepare the retirement of the 3D streaming

Sites / Services round table:

- Gonzalo/PIC: ntr
- Lisa/FNAL: ntr
- Roger/NDGF: ntr
- Tiju/RAL: ntr
- Jhen-Wei/ASGC: ntr
- Ron/NLT1: ATLAS issues seem solved now, but we'll keep a close eye on it
 - ◆ [Maarten: can you confirm that the field type mismatch has been there forever? Ron: yes. Maarten: so probably this was hidden and only appeared now because the usage has increased? Ron: yes this is what we think. Note that our issue is specific to the ACL tables. The issue previously observed at NDGF was a different one related to Chimera.]
 - ◆ [Simone: who will take care of broadcasting this information? Ron: I will broadcast an email to the SRM mailing list.]
- Rolf/IN2P3: will deploy a new version of batch next Monday, will be at risk 12-13 even if it should be transparent.
- Rob/OSG: ntr
- Dimitri/KIT: tomorrow is a holiday in Germany and no one will attend.

- Jan/Storage:
 - ◆ The LHCb issue this morning was related to the intervention yesterday.
 - ◆ Followup on the EOS AATLAS issue from Monday: removing duplicates, only 11 files were affected, most of them related to previous incidents.
- Luca/Databases: ntr
- David/Dashboard: ntr

AOB: MariaDZ, in collaboration with CERN Technical Training is collecting material to shortlist candidate hadoop instructors for a course to take place at CERN. Could experiments please email Maria on:

1. how many people would attend such a course.
2. what level should it be? Introductory or advanced? What time proportion of theory vs hands-on exercises?
3. when would people like the 1st session to take place? We prefer mid-September as the CERN training centre will be unavailable 15/7-25/8 for system upgrades and other maintenance stuff.

Thursday

Attendance: local(Simone, Maria, Jamie, Elisa, Eva, David, Jan, Maarten, MariaDZ, Ignacio);remote(Michael, Rolf, Ronald, Kyle, Jhen-Wei, John, Roger, Lisa, Ian).

Experiments round table:

- ATLAS reports -
- T0
 - ◆ Starting at approx 18:15 yesterday, data transfer attempts to/from T0ATLAS at CERN have been failing massively. Writing gives "Device or resource busy", retrieving data times out on the client side, after 15min of no response. Alarm ticket GGUS:82908 issued. The issue has

been created by the "tilebeam" ATLAS poweruser issuing disk-2-disk copies from t0atlas to the calibration pool. While this is a legitimate use case, it should be used sparsely and not for more than a few tens of files, coordinated with the ATLAS T0 (obviously this was not done). The CERN CASTOR ops killed the transfers under ATLAS request and banned the user. This brought immediately the situation back to normal. The user has been instructed/tortured (I am pretty sure he will not do it again). The ATLAS T0 people will communicate directly with CASTOR ops about the un-banning. Apologies for this mess.

- ◆ There was a problem in a network component (a switch AFAIK) in the CERN CS during the night. Three machines in ATLAS computing Central Services were affected, went in "no_contact" and have been rebooted by the CS operator. The affected machines were part of load balanced services, so the services (Panda/DDM) have been degraded but no issue was observed.
 - T1
 - ◆ NTR (i.e. happy with SARA now; will discuss tomorrow about putting back into export)
-

- CMS reports -
 - LHC machine / CMS detector
 - ◆ Physics ongoing.
 - CERN / central services and T0
 - ◆ One of the Tier-0 boxes rebooted overnight. Set off alarm, responded, and system is recovering
 - ◆ Lost 2 disk servers in EOS and hence lost calibration file. Will modify system so that 3 copies of this particular file are created
 - Tier-1/2:
 - ◆ Regular updates on HC failures - need to investigate further
-

- ALICE reports [↗](#) -
 - ◆ NTR
-

- LHCb reports -
- Users analysis and prompt reconstruction and stripping at T1s ongoing
- MC production at Tiers2

• **New GGUS (or RT) tickets**

- T0:
 - ◆ CERN:
 - T1:
 - ◆ Set Inactive all the CONDDDB access in order to prepare the retirement of the 3D streaming
 - ◆ IN2P3: ask to suspend jobs during the schedule intervention on Monday for GE upgrade [Rolf - this is an AT RISK hence normally transparent. Submission of jobs could be delayed for some minutes hence no reason to either drain or suspend jobs]
-

Sites / Services round table:

- BNL - ntr
- IN2P3 - nta
- NL-T1 - ntr
- ASGC - ntr
- RAL - deployed transfer manager for castor atlas today - went according to schedule
- NDGF - ntr

- FNAL - ntr
- OSG - ntr
- CERN DB - intervention on CMS integration DB in pit to apply patch to workaround network issues found in last TS.
- CERN Grid services - continue working on problems with bad batch performance. Moments of crisis coincide with link to master saturating. Working with programs that interface CEs to batch as they represent a large fraction of the load. Query just for grid queues and q by q. Seems to be a big improvement - this morning job rate was rising but load ok.

AOB: (MariaDZ) Opened GGUS:82907 [↗](#) about the lack of VO value on TEAM tickets reported by ATLAS yesterday. Due to 4 days of german holiday we shall have to wait... Also opened a TEAM ticket to test GGUS:82934 [↗](#) where the VO appears just fine automatically. Conclusion, so far. the problem cannot be reproduced. [Maarten - problem occurred again today to ATLAS]

Friday

Attendance: local(Ian, Jamie, Eric, Simone, David, Jan, Maarten, Ignacio);remote(Michael, Rolf, Kyle, Jhen-Wei, Roger, Gareth, Gonzalo, Onno, Lorenzo, Burt).

Experiments round table:

- ATLAS reports -
- CENTRAL SERVICES
 - ◆ Yesterday at approx 5:30 PM all panda server frontends (6 hosts) went in high load (almost 100% CPU consumption). A user was aggressively polling (every 10 secs or so) all his jobs in Panda. The user has been contacted and "educated". The load went down right after he reduced the frequency. There was no side effect observed in Panda per se. We will put a protection against this in Panda. It is not only LSF or CASTOR @ CERN being hammered...
- T0
 - ◆ NTR
- T1
 - ◆ From RAL: "There is a file transfer problem that appears to be centred in the RAL FTS. We had initially thought there was a Castor/SRM problem, but on investigation that is not the case. Our investigations are ongoing."
 - ◆ 50% jobs for a specific task are failing in INFN-T1 trying to access condition data. Not necessarily a site problem, experts are looking into it.
 - ◆ "Installation" of release 17.2.3.3.1-atlasphysics failed at BNL_CVMFS and in other US CVMFS sites for some reason. Under investigation. Michael - s/w is tagged now so problem fixed.

- CMS reports -
- LHC machine / CMS detector
 - ◆ Physics ongoing.
- CERN / central services and T0
 - ◆ NTR
- Tier-1/2:
 - ◆ Tier-1s: CMS is starting a targeted data reprocessing for ICHEP. This combined with time critical simulation reconstruction.
 - ◆ Tier-2s: We started issuing tickets to look at glexec configuration. We'll switch the tests to indicate a warning if glexec is not working soon.

- ALICE reports [↗](#) -
 - ◆ NTR
-

- LHCb reports - N/A
-

Sites / Services round table:

- BNL - ntr
- IN2P3 - ntr
- ASGC - ntr
- NDGF - sched downtime on Monday for storage between 15:00 and 20:00 Might affect availability of ATLAS and CMS files...
- RAL - file transfer problem resolved 1-2 h ago; believe working ok now; traced to FTS agenda. ANother problem from CMS to do with DNS lookups for addresses at FNAL. Traced to DNS-SEC. Resolved - change made at our end to fix it. 1 d/s out for ATLAS; will follow up. Ext 13 June scheduled outage for 2.1.11-9 update. Long morning or short working day(?)
- PIC - ntr
- NL-T1 - SARA has a warning for maintenance on 1 of 2 tape libs. Some tape fiels can't be staged. Emergency as cartridge insertion was broken and cleaning cartridges could not be inserted. Small performance issue on SARA SRM - now in SRM DB and not N/S DB. Maybe didn't notice before due to that. Causes some errors on some FTS transfers. Rescheduled some cron jobs that were close together. Scheduled down for SARA on Monday 18th for network maintenance.
- CNAF - same probem in FTS for LHCb and CMS VOs. Long standing FTS proxy problem.
- FNAL - ntr
- OSG - ntr

- CERN - ntr

AOB: MariaDZ opened Savannah:129256 [↗](#) , as the intermittent presence of the VO value on tickets opened by ATLAS shifters takes time to be explained. If other experiments have anonymous shifters accounts or ever see no VO value on their tickets, please paste ticket numbers in this savannah ticket or bring them to this meeting. Maarten - one of ATLAS tickets involved was marked "solved" without any comment. What was problem and how fixed?

-- JamieShiers - 22-May-2012

This topic: LCG > WLCGDailyMeetingsWeek120604

Topic revision: r23 - 2012-06-08 - 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