

Table of Contents

Week of 090817.....	1
WLCG Service Incidents, Interventions and Availability.....	1
GGUS section.....	1
Daily WLCG Operations Call details.....	1
Monday:.....	1
Tuesday:.....	3
Wednesday.....	4
Thursday.....	5
Friday.....	6

Week of 090817

WLCG Service Incidents, Interventions and Availability

VO Summaries of Site Availability				SIRs & Broadcasts	
ALICE	ATLAS	CMS	LHCb	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 and select "VO" - this will give all tickets, including team & alarm
- LHCOPN Tickets in GGUS: see 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.
- 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

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

Additional Material:

STEP09	ATLAS	ATLAS logbook	CMS WLCG Blogs
------------------------	-----------------------	-------------------------------	--------------------------------

Monday:

Attendance: local(Jamie, Olof (chair), Gang, Edoardo, Jan, Patricia, Julia, Maria, Diana, Harry, Lola);remote(Angela, Xavier, Jeremy, Gareth, Alexei, Brian, Alessandro (CNAF)).

Experiments round table:

- ATLAS - (Alexei) processing an urgent request for 2TB data replication. Two of the data sets have their unique copies on RAL: when will RAL be back? (RAL report below: *hopefully end of the afternoon*).
- CMS reports -
- ALICE - (Patricia) continuous changes to production versions of AliROOT during the Weekend and as a result there are not many jobs in system but sites shouldn't worry. One of the CREAM CEs back at CERN (ce202) is back since Friday and ticket is closed. Still one of the CREAM CEs (ce201) is down. ALICE is starting migration of VOBoxes to SLC5 at CERN to give feedback (in particular for

the UI) to the grid deployment team.

- LHCb reports -

Sites / Services round table:

- RAL (Gareth) -
 - ◆ Current situation: are in process of restarting. Air-conditioning is back and the cause has (hopefully) been fully understood.
 - ◆ CASTOR is still down in GOCDDB but the service is almost completely back and declared up in GOCDDB hopefully shortly after the meeting.
 - ◆ Tape robot: condensation water dripping into tape report - fortunately the situation wasn't as bad as it was featured and no damage to media or drive. Robot is back and will be accessible once CASTOR is fully back.
 - ◆ Scheduled outage next Thursday for moving ATLAS LFC from non-ATLAS.
 - ◆ (by email):

The RAL Tier1 (RAL-LCG2) carried out an emergency power down following air conditioning failure during the night Tuesday-Wednesday 11-12 August, this was the second event in 2 days. All batch and CASTOR services had to be halted (and remain down) other critical services such as RGMA, the LFC and FTS have remained up the whole time. The failure was caused by the chillers for the new R89 machine room stopping after high pressure triggered coolant flow to be switched off. The cooling system was successfully restarted by 16:00 on Wednesday, 12th however the root cause of the problem has not been identified (although a faulty component is suspected) and the engineers are not yet sufficiently confident that they can assure us we will not suffer a repetition over the weekend. We are also assessing the state of the disk servers, powering them up rack by rack (to keep the load low) - so far the indications are good that no damage has occurred beyond what is normal in a power off situation.

Unrelated to the above problems, on Thursday 13th we also suffered a water leak (condensation) onto our main CASTOR tape robot. STK engineers are on-site assessing the impact of this leak, initial indications are encouraging (as the volumes of water were quite low) however we are very concerned to ensure no contamination of the media and checks are ongoing.

After consultation with our UK experiment contacts we have decided that the risk to equipment of another forced power down is too high to justify restarting the service before the weekend. Work on understanding of the problem is making good progress and we believe it is likely (80%) that by Monday, we will be able to begin the restart. Once we have sufficient information from the engineers we will decide how the start-up will be managed, we have not yet decided if we will immediately start at full capacity or gradually ramp back up.

We appreciate the impact this break in service has on the LHC experiments but do not wish to risk further emergency power downs unless absolutely necessary. If any experiment has urgent requirements that make restart of service vital we would appreciate their input as soon as possible so we can take that into account as we balance startup risks against experiment requirements.

- FZK (Xavier) - about 1.5 hours ago several dCache pools disabled themselves because of a glitch in GPFS.
- CNAF (Alessandro) - NTR
- GridPP (Jeremy) - NTR
- ASGC (Gang) - Degradation in CASTOR service (both ATLAS and CMS). Local operation team is

Monday:

investigating.

- CERN

- ◆ (Edoardo) Networking problem last week: DHCP broadcast overloaded a router CPU. The originating IP service was disconnected, after which the router situation went back to normal. Networking team is in contact with owner of the machine that caused the storm but the cause has not yet been understood.

AOB:

Tuesday:

Attendance: local(Harry, Simone, Diana, Patricia, Jamie, Julia, Gang, Maria, Olof (chair));remote(Xavier, John Kelly (RAL), Michael, Brian, Alessandro).

Experiments round table:

- ATLAS - (Simone) reprocessing: postponed until August 31st because of internal reasons. Site related issues - see the RAL report below. Also, noticed some 40% failures transferring data into ASGC. The error returned is 'Unknown error'.
- CMS reports -
- ALICE - (Patricia) VOBox issue with firewall definition at CERN. Problem had been reported from several sites where ALICE VOBoxes couldn't connect to the Alien DB at CERN. The port, 8084, had been closed recently and must now be opened. Patricia said that the port has been open in the past (since ever) and it's unclear why it was closed. Ongoing investigations together with the networking team. CERN CREAM CEs: ce202 is ok and working fine. ce201 is being restored in a similar manner.
- LHCb reports -

Sites / Services round table:

- FZK (Xavier) - GPFS glitch reported yesterday: had to fix some side effects this morning, which had not been noticed after the glitch yesterday. The effect was that some dCache pools were down (reporting IO errors). Everything is back to normal now.
- RAL (John) - batch system is back in full productions including all worker nodes since 09:30 this morning. Still some SAM tests failures. Lost a disk server and there might be some ATLAS files lost with it. Brian: *yes, we did lose 99k files out of 4M of the MCDISK data. A list of lost files has been produced. Currently in process of removing the files from the catalogs. There may be some job failures in the coming days.* Because of the kernel security issue RAL has disabled interactive logins for outsiders to RAL UIs. Harry: *are your worker nodes back?* yes, 09:30 this morning.
- BNL (Michael): two small things - all interactive nodes have been patched for preventing effects from the security vulnerability. Also, working on the site name consolidation. Today, we're doing the first phase and plan to complete the name change to within a week from now. The availability of relevant experts (FTS, ...) has been confirmed.
- CNAF (Alessandro) - nothing special to report except that like for other sites CNAF has been patching the frontend nodes for the security problem.
- ASGC (Gang) - ATLAS seems to be more affected than CMS from the CASTOR SRM degradation. Have not received any update from local experts since yesterday. One reason why ATLAS is more affected than CMS may be because they use space tokens...?

AOB:

Tuesday:

Wednesday

Attendance: local(Jamie, Simone, Roberto, Patricia, Gang, Lola, Oliver, Eva, David, Maria, Diana, Olof (chair));remote(Gonzalo, Angela, Tiju Idiculla (RAL), Alessandro).

Experiments round table:

- ATLAS – (Simone) One issue with PIC starting from this morning: trying to read data from PIC results in SRM source problem and write into PIC gives a strange gridftp error. A ticket has been sent. David: *there is a OPN GGUS ticket about this problem and it's related to some instability in the primary link. The performance dropped to 1Gbit/s from 5 yesterday. Currently running on the backup link.* The other small point is that RAL confirmed that they finished the cleanup of the diskserver that was lost two days ago. The catalogues are now consistent locally and the DQ2 catalogues will be done today.
- CMS reports -
- ALICE – (Patricia) Update on the CERN firewall problem for Alien DB mentioned yesterday: the request to open 8084 port here at CERN has been done and a security scan will be run today after which the port will be opened to outside. A workaround has been put in place by ALICE in the meanwhile.
- LHCb reports - (Roberto) Currently very little activity (~1kjobs in system). In the meanwhile a cleanup of MC data no longer used by the community. Few issues: SARA space token for master MC data is running out of space. Concerning the production, LHCb is running only on half of the Tier-1s for various reasons.
 - ◆ NL: SARA down time
 - ◆ RAL: now back in action
 - ◆ IN2P3: some problems where jobs are failing due to a memory limitation. GGUS ticket open
 - ◆ CNAF: sqlite issue. Seems to be fixed but from time to time some WN fails. Seems to be some cron related job.

Sites / Services round table:

- PIC (Gonzalo) – nothing special. Just to confirm the networking instabilities reported by ATLAS
- FZK (Angela) – another GPFS glitch. Replaced one blade in one of the routers.
- RAL (Tiju) – NTR
- CNAF (Alessandro) –
 - ◆ sqlite problem reported by LHCb above: the problematic nodes have been identified and isolated.
 - ◆ Progress concerning ATLAS problem on memory allocation in LSF jobs: 64bit WNs have now been allocated for the ATLAS queues. This solves the problem seen with memory allocation limit on the 32bit nodes. Simone: *ATLAS will confirm. Probably best if somebody from CNAF can participate at the ATLAS operation meeting tomorrow at 15:30.* Will do.
- ASGC (Gang) – CASTOR SRM degradation was fixed late yesterday. However, this morning the migration of CMS files stopped. Local administrators didn't find any problem and will request help from the CASTOR developers.
- CERN (Olof) - NTR

AOB:

Thursday

Attendance: local(Jamie, Julia, Simone, Miguel, Lola, Jacek, Roberto, Diana, Gang, Maria, Olof (chair));remote(Gareth, Angela, Gonzalo, Brian, Xavier, Andreas).

Experiments round table:

- ATLAS - (Simone) Two days ago 170TB was deployed MCDISK area at CERN. Allows to replicate all merged AOD data and this activity started today. It's the first time we start a large import to CERN. Started at 2GB/s, which is very good, and now we are at the tail with 600-700MB/s. PIC connectivity problem: connectivity between PIC and CNAF doesn't seem to work, even over the backup route. Gonzalo: *will look into this problem with the network problem. In principle the primary OPN link is up since 11am.* Simone: saw errors in the last hours for transferring data from PIC.
- CMS reports -
- ALICE -
- LHCb reports - (Roberto) Not much to report. Usual MC production ongoing, should be finished by tomorrow but new physics requests are coming. Finished massive cleanup of data reported yesterday. This freed 230TB of storage capacity. Question for CASTOR @ CERN concerning Philippe's request for reshuffling: LHCBDATA low in space and request for a new pool for histograms. Miguel: *the reason why we didn't answer yet is because we are cleaning up out-of-warranty diskservers. This should finish by next Tuesday. Thereafter we have some spare capacity for fulfilling the new space requirements and in principle we would not need to remove capacity from LHCRAW.* CNAF is setting up a new StorM endpoint and initial tests are ok.

Sites / Services round table:

- RAL (Gareth) - LFC outage this morning when the ATLAS LFC was split. There was an overrun on the announced outage in GOCDB but the work completed successfully but the failing SAM tests are still failing because pointing to the wrong LFC, which will be fixed. Brian: *wondering whether or not if it is possible to get a unified approach for the process of reporting and what needs to be done when sites temporarily loose diskservers?* Maria: *this was discussed months ago.* Jamie: *the daily operations meeting is the place where you can raise it but not the place where we will study it. There should be a separate working group on the topic (probably Andrea Sciaba's storage WG?).* Gareth: *will mail wlcg-scod list with the RAL input.*
- FZK (Angela) - next week FZK will move SL4 nodes to SL5. Will start with half of them on Monday, draining the nodes from running jobs. At the same time new CEs will put in pointing to the SL5 capacity. Thereafter the second half of the WNs will be processed in the end of the week.
- PIC (Gonzalo) - will follow up the networking issue. A different issue: next week we have a scheduled several hours downtime (monthly) at PIC. Up to now they normally inserted the scheduled downtime in GOCDB 7 days ago. It used to be possible to broadcast a warning immediately but this doesn't seem to be possible anymore...? Jamie: *I think it should be put back.* Maria: *there was a message by Maite about simplified broadcast scheme some time ago.* Gareth: *we also noticed the change where the GOCDB interface was simplified where the broadcasting is automated through some rules.* Miguel: *there used to be a tick-box, which has been removed in the new version.* [After the meeting Maria found the message mentioned above. Some material at <http://indico.cern.ch/getFile.py/access?sessionId=5&resId=0&materialId=0&confId=64396> It is still unclear if this explains the problem seen by Gonzalo.]
- CNAF (Alessandro) - NTR
- ASGC (Gang) - tape migration problem reported yesterday has been fixed this morning.
- CERN (Miguel) - on Monday we will start a new linux software upgrade, freeze. CASTORALICE will be upgraded to 2.1.8-10.

AOB:

- MariaDZ: USAG meeting in a week. Agenda on <http://indico.cern.ch/conferenceDisplay.py?confId=66479>. Experiments, please, select the most recent report from https://gus.fzk.de/pages/metrics/download_escalation_reports_vo.php and clean up your tickets.
- Jamie: alarm ticket tests, would it be feasible to do one before the next GDB? Maria: it has to be done and completed the week before the GDB. A test would be possible and useful but Gunter, who normally orchestrate this, leaves on holidays next week. Jamie: then a test should be done before the October GDB, despite the EGEE conf in between.

(MariaDZ) Plan done in

https://twiki.cern.ch/twiki/bin/view/EGEE/SA1_USAG#Periodic_ALARM_ticket_testing_ru. Result report will be attached to <https://savannah.cern.ch/support/?109566>

Friday

Attendance: local(Simone, Lola, Gang, Roberto, Patricia, Alessandro, Olof (chair));remote(Gareth, Xavier, Gonzalo, Dantong, Tiju, Brian, Alessandro).

Experiments round table:

- ATLAS - (Simone) No problems with sites or services to report. Follow-ups from previous days: 1) PIC network problems and the subsequent transfer problems between PIC-CNAF. Can confirm that the transfer channels, including PIC-CNAF, were ok after the primary route was back. It is also understood why PIC-CNAF wasn't working on the secondary route and the problem will be fixed soon. 2) it was reported a few days ago that all CNAF WNs were moved to 64bit hardware in order to solve an issue with ATLAS jobs being killed due to memory limit. After having checked yesterday, Simone can confirm that the problem has indeed been solved. 3) Concerning the BNL site name consolidation: yesterday there was a discussion in the ATLAS operation meeting. BNL is now ready to move to 'BNL-ATLAS', which will require an coordinated action for all Tier0/1 FTS servers to change the site-name. Dantong: *BNL-ATLAS will be the permanent name. This involves all Tier-0 and Tier-1 administrators following Gavin's procedure. There is one important update to the original procedure required: update the 'site name' to the actual one (BNL-ATLAS) in the procedure. The proposal is to do this on Tuesday 25/8/2009 at 9am EDT, 3pm CEST. Also all the ATLAS Tier-1 FTS admins must apply the procedure* Simone: propose to look after the update at 3pm, and follow-up in case there seem to be any of them having problems. Gonzalo: *the wlcg-tier1-contacts@cernNOSPAMPLEASE.ch mailing can be used for coordinating this operation.* Agreed. Alessandro: *would be good to know the exact date/time for when the Tier-1s should change FTS.* Dantong: propose that Tier-1s does the update 15mins after the start time, i.e. 15:15 on Tuesday the 25th.
- CMS reports -
- ALICE - (Patricia) announcement for all the sites: after the latest upgrade to 2.17 of Alien a number of LDAP change are required. This is being done site per site. For the moment the job information is wrong. This update will take the rest of today. Going over each of the individual sites was useful exercise that revealed some problematic Tier-2: CYFRONET-LCG2 (see GGUS 51064, https://gus.fzk.de/ws/ticket_info.php?ticket=51064&from=ID) and IN2P3-LPSC (GGUS 51065, https://gus.fzk.de/ws/ticket_info.php?ticket=51065&from=ID).
- LHCb reports - (Roberto) Huge activity ongoing: many (7000) of MC jobs running in the system + other jobs from distributed analysis. The active testing activity for DIRAC SL5 found a working combination of packages for building a working version. Testing is ongoing and there is good hope

there will soon be a SL5 certified version of DIRAC.

Sites / Services round table:

- PIC (Gonzalo) - NTR
- BNL (Dantong) - network glitch yesterday (4-5mins) between WNs and CEs. Affected some production jobs.
- RAL (Gareth) - NTR
- CNAF (Alessandro) - NTR
- ASGC (Gang) - migrations for CMS stopped today again. Local administrator is following up.
- CERN (Olof) - NTR
- FZK (Xavier) - announced downtime on Tuesday 25th 10:00 - 10:30 for upgrading dCache SRM service. *After the meeting Xavier submitted a correction stating that the ce3, ce-cream1-fzk upgrade announced in the meeting was a misunderstanding: the CEs will not be upgraded to SL5 but only reconfigured to point towards workernodes with SL5 (as was already announced yesterday) . The two CEs will have a downtime for 2 hours (2 hours not available in the information system). This is only used to ensure the right information in the TopLevel BDII's for right job matching. No jobs will be lost.*

AOB: *

This topic: LCG > WLCGDailyMeetingsWeek090817

Topic revision: r11 - 2009-08-21 - unknown



Copyright &© 2008-2020 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