

# Table of Contents

<b>Week of 130617.....</b>	<b>1</b>
Daily WLCG Operations Call details.....	1
WLCG Availability, Service Incidents, Broadcasts, Operations Web.....	1
General Information.....	1
Monday.....	1
Thursday.....	3

# Week of 130617

## 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](#)

The scod rota for the next few weeks is at [ScodRota](#)

## WLCG Availability, Service Incidents, Broadcasts, Operations Web

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

## General Information

General Information			GGUS Information	LHC Machine Information
<a href="#">CERN IT status board</a>	<a href="#">WLCG Baseline Versions</a>	<a href="#">WLCG Blogs</a>	<a href="#">GgusInformation</a>	<a href="#">Sharepoint site - LHC Page 1</a>

## Monday

Attendance:

- local: AndreaS, Raja (LHCb), Andrew (LHCb), Luca, Thomas, MariaG, Ivan, Stefan, Maarten, MariaD, Marcin
- remote: Luc (ATLAS), Tommaso (CMS), Xavier (KIT), Pepe (PIC), Lisa (FNAL), Rolf (IN2P3), Boris (NDGF), Tiju (RAL), Onno (NL-T1), Wei-Jen (ASGC), Rob (OSG), Paolo (CNAF)

Experiments round table:

- ATLAS reports (raw view) -
  - ◆ T0/Central services
    - ◇ NTR
  - ◆ T1
    - ◇ SARA failing transfers, "RQueued" Transfer Errors. [GGUS:94899](#). Understood (DNS problem with tape library system)
    - ◇ SARA SDT: dCache databases to be upgraded from Postgresql 8.4 to 9.2, and moved to new hardware.
- CMS reports (raw view) -
  - ◆ Nearly uneventful week, which is a good sign! Not a lot to report
  - ◆ Production + 2011 reprocessing ongoing
  - ◆ Any update on general networking connectivity for Russia? We also saw problems with our sites, and we have a Savannah on this: [SAV:138162](#). Unfortunately the answer up to now

has been that they do not know the timescale for recovery...

- ◆ Two PhEDEx instances gown down at T1\_CNAF in the night due to VM problems, fixed early morning.

Maarten: also ALICE had to stop using some Russian T2s. There is a ticket open (GGUS:94540 [↗](#), mentioned in the minutes of past Monday), other experiments affected should follow and contribute to it. There is no estimate of when the link will be fully operational. So far, LHCb and ATLAS are not affected.

- ALICE -
  - ◆ KIT: Xrootd configuration error was corrected already on Wed and during the last days the job rates have been ramped up successfully again.
- LHCb reports (raw view) -
  - ◆ Incremental stripping campaign in progress and MC productions ongoing
  - ◆ T0:
  - ◆ T1:
    - ◇ GridKa : IDLE pilots preventing new jobs starting (GGUS:94898 [↗](#)). Ticket opened yesterday and escalated to ALARM this morning. Solved now - possibly due to a black hole node.
    - ◇ GridKa : Pilots going through GridKa brokers fail due to sandbox being lost (GGUS:94462 [↗](#))
    - ◇ GridKa : Jobs at GridKa failing due to local directory problems (GGUS:94471 [↗](#)). This affects many jobs at GridKa, and merging jobs with large numbers of input files have very low success rate.
  - ◆ Other: Web Server overloading continues (GGUS:94824 [↗](#)) - possible problem now due to the afs server underneath. Under investigation and affects jobs at many sites, some more than others. [Stefan: the web server has been duplicated, so it's not overloaded anymore: the excessive load falls entirely on AFS and still needs to be addressed]

Sites / Services round table:

- ASGC: ntr
- CNAF: ntr
- FNAL: ntr
- IN2P3: after migrating 1/4 of our WNs to SL6 we notice a big increase in the amount of memory used by jobs. Not clear yet if this is specific to certain jobs or VOs, still investigating. Maarten: many sites use SL6 in production since months and no similar problems were reported so far. Rolf: some parameters can be tuned to improve, also at compilation time. Luca suggests to contact the CERN batch experts to see if they have encountered this problem.
- KIT: ntr
- NDGF: ntr
- NL-T1: ongoing downtime to upgrade the dCache database, it should end around 1830 CEST.
- PIC: ntr
- RAL: tomorrow there will be a downtime for electrical maintenance.
- OSG:
  - ◆ during the weekend ~10 hours of connection problems to the CERN BDII; it was due to a misbehaving node at CERN which was then taken out of the alias
  - ◆ noticed problems in the propagation of tickets from GGUS to the OSG messaging system, which may have affected tickets during the whole weekend; in contact with the GGUS developers to understand the cause
- CERN storage services:
  - ◆ next Monday we would like to have a downtime from 1600 to 1800 CEST to reboot all CASTOR nodes; it would overlap with a CASTOR database intervention to fix a bug and which would cause all activity to be suspended for 30'. The CASTOR team will send an

announcement today and if no objections are made the intervention will be scheduled in the GOCDDB.

- ◆ The EOS cluster is undergoing a reboot campaign from today for a week; it will be totally transparent.
- Dashboards: ntr
- Databases: ntr
  
- GGUS:
  - ◆ ATLAS please take care of tickets assigned to you, namely GGUS:90428[?](#), GGUS:94294[?](#), GGUS:94784[?](#).
  - ◆ Slides and drills for the WLCG MB attached to this page end.

AOB:

## Thursday

Attendance:

- local: Luc/ATLAS, Raja/LHCb, Andrew/LHCb, Ivan, Gavin, Maarten, MariaD
- remote: Saerda/NDGF, Michael/BNL, Xavier/KIT, John/RAL, Lisa/FNAL, Ronald/NL-T1, Marc/IN2P3, Jose'/CMS, Pepe/PIC, Wei-Jen/ASGC

Experiments round table:

- ATLAS reports (raw view) -
  - ◆ T0/Central services
    - ◇ CERN-PROD: jobs failing with stage-out errors GGUS:94978[?](#). internal issues (stuck LDAP lookups) but also extra-load from ATLAS (Sim@P1). Ongoing
    - ◇ Pending ggus tickets closed
  - ◆ T1
    - ◇ SARA-MATRIX no restart after DT (i.e. no pilots entering). GGUS:94952[?](#). Pb with a specific WN (cvmfs directories remounted). A priori OK
    - ◇ NDGF-T1 SRM errors GGUS:94979[?](#). Ongoing (dCache issue)
    - ◇ RAL lost file GGUS:94969[?](#). Cannot be recovered
    - ◇ RRC-KI transfer failures from all clouds GGUS:94827[?](#). Solved ( border router losing routes to CERN via LHC OPN)
- CMS reports (raw view) -
  - ◆ Quiet since Monday. Just a couple of minor issues with CNAF and RAL
    - ◇ Black hole WN at CNAF taken out for investigation GGUS:94958[?](#)
    - ◇ Some files not accesible at RAL GGUS:94959[?](#)
- ALICE -
  - ◆ KIT: concurrent jobs cap had to be lowered again to 3k yesterday evening; now back at 4k, will try for higher values later today.
    - ◇ The nominal level would be O(10k). Unfortunately the cause of the underlying problem has eluded the experts so far.
  - ◆ NIKHEF: upgraded to WLCG VOBOX and the VM got more memory added as well, thanks!
- LHCb reports (raw view) -
  - ◆ Incremental stripping campaign in progress and MC productions ongoing
  - ◆ T0:
  - ◆ T1:

- ◇ GridKa : Pilots going through GridKa brokers fail due to sandbox being lost (GGUS:94462 [↗](#))
- ◇ GridKa : Jobs at GridKa failing due to local directory problems (GGUS:94471 [↗](#)). This affects many jobs at GridKa, and merging jobs with large numbers of input files have very low success rate.
- ◇ RAL : Problem starting jobs. Seems to have gone away for now (starting this morning). Internal tickets opened.
- ◆ Other: Web Server overloading significantly alleviated with 3rd web frontend (GGUS:94824 [↗](#)). Affects jobs at many sites, some more than others.

Sites / Services round table:

- ASGC: ntr
- BNL: ntr
- FNAL: ntr
- IN2P3: ntr
- KIT:
  - ◆ we found a broken tape and nine CMS files were lost; CMS has been informed
  - ◆ a server doing most of the staging for LHCb suffered file system problems and staging has been moved to another machine, while the old one is being analysed to understand the cause
  - ◆ will ask our experts to comment on the tickets opened by LHCb
- NDGF: the recent dCache upgrade introduced an instability, still under investigation
- NL-T1: ntr
- PIC: ntr
- RAL: the recent intervention on the electrical system was completed without issues. Another intervention, also expected to have no impact, has been scheduled for the 25th as "at risk"
- CERN batch and Grid services: after some recent problems, the SLC6 capacity is now almost fully available and it accounts to ~20% of the full capacity of LXBATCH. A further increase (and a reduction of the SLC5 capacity) will start at the end of summer
- Dashboards: ntr
- Databases: two interventions on the physics databases have been scheduled for next Tuesday and Wednesday, both from 0900 and expected to take ~30'. During that time no I/O will be possible. The full list of databases affected by each intervention is available in the ITSSB (Tuesday [↗](#), Wednesday [↗](#))
- GGUS: The Footprints/GGUS trouble of last weekend 15-16 June, reported by Rob (OSG) last Monday was due to a coming change in GGUS that was put by mistake, also on the production system. Guenter (GGUS) rapidly backtracked and Soichi (OSG) is getting ready for the official release on 10 July. Details in Savannah:138191#comment2 [↗](#)

AOB:

---

This topic: LCG > WLCGDailyMeetingsWeek130617

Topic revision: r15 - 2013-06-20 - MaartenLitmaath



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