

# Table of Contents

<b>Week of 110808.....</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.....	4
Thursday.....	6
Friday.....	8

# Week of 110808

## 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 (AndreaV, Elena, Jarka, Jamie, Tom, Jose, Mike, Ulrich, Jan, Steve, Alessandro, Eva, Nilo, Lola); remote (Michael, Giovanni, Jhen-Wei, Tiju, Ron, Catalin, Thomas, Pepe, Dimitri, Rob; Daniele, Vladimir).

Experiments round table:

- ATLAS reports -
  - ◆ CERN-PROD: ATLAS TAG DB (Archive DB) flipping from "green" to "red" status on SLS monitor. First observed around 6pm on Sunday. Is something wrong going on? [Jarka: ATLAS DBAs are investigating and will contact the IT-DB DBAs.]
  - ◆ INFN-T1: SRM down on Saturday. GGUS:73236 TEAM escalated to ALARM. The problem was on the StoRM BE which needed a power cycle in order to reset the system.
    - ◇ It seems GGUS did not send properly ALARM notification. Ticket was submitted by /atlas/team member, who is not /atlas/alarm-er. Escalation by /atlas/alarm-er apparently did not reach SMS system of the site. [AndreaV: is this similar to the issues observed in previous weeks? Jarka: yes, it's the third time this happens with GGUS, will contact the GGUS experts to discuss this. Alessandro: similar pattern though slightly different in the details, and it always happened during the weekend.]
  - ◆ [Steve: reminder, there will be an intervention on ATLAS cvmfs at 4pm CERN time].

- CMS reports -
  - ◆ LHC / CMS detector
    - ◇ Fills for physics
  - ◆ CERN / central services
    - ◇ Still monitoring glitches in SLS: INC:056362 [↗](#) . Easy to identify by computing shifters since not only the available disk space at the castor pools drops to zero but also the total capacity. [Jan: would tend to work on this with standard (not top) priority, this is generally the agreement with CMS for monitoring issues. Jose: OK.]
  - ◆ T1 sites:
    - ◇ Still investigating issue of job submission to KIT with ProdAgent (Savannah:122611 [↗](#))
    - ◇ The data reprocessing ran over the weekend with no problems.
- ALICE reports [↗](#) -
  - ◆ General information: We are experiencing again connection problems with lgvoms.cern.ch . It is a known instability of the ALICE VOMS servers. When will be deployed the upgrade that solves this problem? [Ulrich: will forward the question to Steve.]
  - ◆ T0 site
    - ◇ Nothing to report
  - ◆ T1 sites
    - ◇ Nothing to report
  - ◆ T2 sites
    - ◇ Nothing to report
- LHCb reports -
  - ◆ New GGUS (or RT) tickets:
    - ◇ T0: 0
    - ◇ T1: 2
    - ◇ T2: 0
  - ◆ Issues at the sites and services
    - ◇ T0
      - SetupProject timeout under investigation (GGUS:73177 [↗](#))
    - ◇ T1
      - IN2P3: pilots aborted at cccreamceli06 (GGUS:73246 [↗](#))
      - SARA: Jobs failed to upload data at cream2 (GGUS:73244 [↗](#))

Sites / Services round table:

- BML: ntr
- CNAF: two other problems on top of those already mentioned:
  - ◆ one BDII issue solved on Sunday evening, ticket opened with the developers
  - ◆ huge load of logging system, affected Storm, solved today
- ASGC: ntr
- RAL: reminder, site firewall will be rerouted tomorrow (it is in GOCDB)
- NLT1: ntr
- FNAL: ntr
- NDGF: ntr
- PIC: reminder, will upgrade dcache tomorrow
- KIT: ntr
- OSG: ntr
  
- Databases services: ntr
- Grid services: nta
- Storage services [Jan]: would like to complete the move of CASTOR to DNS load balancing this

week, will affect castoratlas, castorcms and castorpublic.

AOB: none

## Tuesday:

Attendance: local (AndreaV, Elena, Jamie, Mke, Eva, Jan); remote (Michael, Xavier, Ronald, Caalin, Giovanni, Marc, Thomas, Jhen-Wei, Jeremy, Tiju, Pepe; Daniele, Vladimir).

Experiments round table:

- ATLAS reports -
  - ◆ CERN-PROD:
    - ◇ We are seeing high load on ATLR. The delay in replication from online to offline DB was observed during the night and in the morning. DB on call was contacted.
    - ◇ The 3rd node has been successfully added to ADCR. Thanks.
  - ◆ The problem with notification of GGUS upgraded to ALARM seen over weekend was investigated by INFN. INFN updated their SMS notification system. Thanks. Test GGUS:73317 [↗](#) opened to follow issue with INFN-T1. We're in touch with GGUS to see what changes in workflow can be done so that ALARM tickets do reach sites on time. [Andrea: was this fixed at KIT or CNAF? Giovanni: this was fixed at CNAF. Elena: in contact also with central GGUS team to prevent this happening at other sites in the future; this only happened for TEAM tickets escalated to ALARM]
- CMS reports -
  - ◆ Daniele Bonacorsi CRC for this week.
  - ◆ LHC / CMS detector
    - ◇ Fills for physics
  - ◆ CERN / central services
    - ◇ NTR
  - ◆ T1 sites:
    - ◇ all T1 sites: the data reprocessing ran over the weekend up to a >99% completion level. Dealing with the tails.
    - ◇ KIT: still investigating the issues with the job submissions (Savannah:122611 [↗](#))
    - ◇ KIT->FNAL: Degrading phedex transfers KIT->FNAL due to too many parallel transfer tries: already fixed (Savannah:122696 [↗](#)). [Xavier: in which way is there a risk to overload KIT? Daniele: as reported yesterday, there are attempts to do many transfers from KIT to FNAL. Catalin: there was a configuration at FNAL in which KIT disks could not cope with the high number of transfer requests - this is now corrected.]
    - ◇ IN2P3: apt-get error while deploying CMSSW\_4\_1\_7\_patch2 at IN2P3 (Savannah:122699 [↗](#))
    - ◇ PIC: scheduled downtimes (network 8am-10am, storage 10am-6pm, virtual servers 10am 6pm)
  - ◆ T2 sites:
    - ◇ recovering from recent issues in some US T2s (thunderstorm+powercut in Nebraska, power cut in UCSD, maintenance in Purdue)
- ALICE reports [↗](#) -
  - ◆ T0 site
    - ◇ Nothing to report
  - ◆ T1 sites
    - ◇ Nothing to report
  - ◆ T2 sites

Tuesday:

◇ Nothing to report

- LHCb reports -
  - ◆ Experiment activities
    - ◇ A lot of data last day.
    - ◇ It was discovered RAWIntegrityAgent stuck, as result data transfer from pit to Castor was blocked. Fixed by restarting agent. FTS transfer from CERN to GRIDKA was blocked starting from weekend due to few wrong requests. No problems after requests were removed. Backlog of not processed files at GRIDKA.
  - ◆ New GGUS (or RT) tickets:
    - ◇ T0: 0
    - ◇ T1: 0
    - ◇ T2: 0
  - ◆ Issues at the sites and services
    - ◇ T0
      - SetupProject timeout under investigation (GGUS:73177 [↗](#))
    - ◇ T1
      - IN2P3: pilots aborted at cccreamceli06 (GGUS:73246 [↗](#))
      - SARA: Jobs failed to upload data at cream2 (GGUS:73244 [↗](#))

Sites / Services round table:

- BNL: ntr
- KIT: ntr
- NLT1: ntr
- FNAL: ntr
- CNAF: ntr
- IN2P3: ntr
- NDGF: ntr
- ASGC: ntr
- Gridpp: ntr
- RAL: intervention has completed successfully
- PIC: ntr
  
- Dashboard services: ntr
- Database services: ntr
- Storage services: reminder of two interventions tomorrow
  - ◆ transparent switch to DNS load balancing for Alice, LHCb, CMS and public CASTOR in the morning
  - ◆ transparent intervention on EOS for ATLAS and CMS in the aftrenoon

AOB: none

## Wednesday

Attendance: local (AndreaV, Elena, Jarka, Ulrich, Jamie, Jan, Edoardo, Mike, Luca, Lola); remote (Micheal, Catalin, Jhne-Wei, Tiju, Thomas, Marc, Giovanni, Ron, Dimitri, Rob; Daniele, Vladimir).

Experiments round table:

- ATLAS reports -
  - ◆ CERN-PROD:
    - ◇ SRM was down: CASTOR nameserver overload (ALARM GGUS:73350),T0 was suffering from the slow response (ALARM GGUS:73355 [↗](#)). Fixed in less than 1 h.

Thanks. [Andrea: did the notification to the piquet work normally, after the recent GGUS issues? Jan: yes, the ALARM notification reached the piquet as expected.]

◇ Pending LSF jobs on ATLAS public share (GGUS:73348). [Ulrich: all pending jobs seem to have been scheduled within 1h, which is quite good and normal when running on a public shared queue. This is a share that can be used if noone else is using it, so it seems to be working according to design.]

· [Elena: there may be problems due to simultaneous access to the COOL database if all these jobs are started simultaneously. Jarka: this is a separate issue from LSF anyway - and we are trying inside ATLAS to have mechanisms to avoid database contention issues.]

◆ T1

◇ INFN-T1: problem with transfers to SCRATCHDISK (GGUS:73348).

◇ TAIWAN-LCG2: power failure in network switch. The problem was noticed by the site and fixed quickly. Thanks.

• CMS reports -

◆ LHC / CMS detector

◇ Fills for physics

◆ CERN / central services

◇ VOMS service low availability observed by the CMS shifter, not a glitch, GGUS:73349 opened, promptly fixed (it was related to voms-admin actions) [Ulrich: yesterday between 16:00 and 18:00 CEST voms-admin for LHC VOs was unavailable as shown in SLS. Cause was sysadmin carelessness.]

◇ SRM-CMS glitches observed by the CMS shifter, IT SSB stated that all CASTOR SRMs actually reported errors, due to an overload on the DB behind the central CASTOR nameserver. Fixed (reduced number of nameserver threads) within the day. No tickets. [Jan: if these glitches are those observed since the beginning of the week, it is unlikely that they would be solved by the recent DNS move. Daniele: thanks, will tell the shifters that this issue may still reappear then.]

◇ CASTORCMS glitches in SLS monitoring, 5 in 12 hrs in the time window 12:00-24:00 Aug 9th, all smooth since the nameserver DB issue was solved

◆ T1 sites:

◇ KIT: in scheduled downtime, tape system maintenance, SE functional and all CMS traffic goes on transparently (using disk buffer and not the tape BE of course)

◇ ASGC: unexpected power failure caused unscheduled down of a SRMv2 link (see here), no major impact on CMS transfers

◆ T2 sites:

◇ Nothing to report

• ALICE reports -

◆ T0 site

◇ Nothing to report

◆ T1 sites

◇ Nothing to report

◆ T2 sites

◇ Nothing to report

• LHCb reports -

◆ Experiment activities:

◇ Data processing and stripping. Backlog of waiting jobs at GRIDKA.

◆ New GGUS (or RT) tickets:

◇ T0: 0

◇ T1: 1

◇ T2: 0

◆ Issues at the sites and services

◇ T0

◇ T1

- GRIDKA: Read only Shared Area for SAM jobs with software installation (GGUS:73343 [↗](#)); Fixed

Sites / Services round table:

- PIC [Pepe]: ntr
- BNL: ntr
- FNAL: ntr
- ASGC: UPS power failure this morning, should be fixed now
- RAL: ntr
- NDGF: one cluster in Oslo is not accepting connections as it had a security incident, which is under investigation
- IN2P3: ntr
- CNAF: ntr
- NLT1: ntr
- KIT: ntr
- OSG: ntr
  
- Grid services: nta
- Database services: ntr
- Network services: ntr
- Dashboard services: ntr
- Storage services [Jan]:
  - ◆ There was an issue caused by the new checksum verification tool, now a fixed version is being installed.
  - ◆ Transparent intervention this morning on Castor DNS has completed without problems.
  - ◆ EOS Interventionis ongoing as we speak, there was a minor glitch which has immediately been fixed.

AOB: none

## Thursday

Attendance: local (AndreaV, Elena, Ulrich, Steve, Jan, Jamie, Mike); remote (Xavier, Ronald, Jhen-Wei, Kyle, Catalin, Thomas, Giovanni, Gareth, Marc; Daniele, Vladimir)

Experiments round table:

- ATLAS reports -
  - ◆ CERN-PROD:
    - ◇ We continue to observe a high load on ATLR offline DB caused by high submission rate of T0 jobs. DBA and T0 experts are dealing with the problem. [Elena: high load is because each job opens/closes several connections.]
  
- CMS reports -
  - ◆ LHC / CMS detector
    - ◇ Fills for physics
  - ◆ CERN / central services
    - ◇ Nothing to report
  - ◆ T1 sites:

- ◇ ASGC-1: submission of MC prod jobs failed this morning. Fixed on the site side (WN misconfiguration). Resubmitting now. [Jhen-Wei: discussed with the colleague in Taiwan, this was due to an issue with the firewall of some WNs, should be fixed now.]
- ◇ ASGC-2: transfer issues to UCSD and Vienna of urgent SUSY samples (Savannah:122785 [↗](#), Savannah:122786 [↗](#)). [Jhen-Wei: this was related to the UPS problem the previous day, should be ok now.]
- ◆ T2 sites:
  - ◇ Nothing to report
- ALICE reports [↗](#) -
  - ◆ T0 site
    - ◇ Connection problems with some CAF nodes due to some change in the LDAP configuration. Problem solved. Related to that issue there is SNOW ticket open INC:INC054828 [↗](#) in order to customize the ldap configuration via Quattor to use AliEn LDAP instead of CERN LDAP. [Steve: will follow up offline with ALICE.]
  - ◆ T1 sites
    - ◇ Nothing to report
  - ◆ T2 sites
    - ◇ Nothing to report
- LHCb reports -
  - ◆ Experiment activities:
    - ◇ Data processing and stripping. Backlog of waiting jobs at GRIDKA.
  - ◆ New GGUS (or RT) tickets:
    - ◇ T0: 0
    - ◇ T1: 0
    - ◇ T2: 0
  - ◆ Issues at the sites and services
    - ◇ T0
      - Problem with CASTOR, all transfers to CERN were blocked (INC:058607 [↗](#)).
      - [Vladimir: there was also a problem with VOMS, no ticket was opened for this.]
    - ◇ T1
      - ntr

Sites / Services round table:

- KIT: at 12:00 local time GPFS crashed, some ATLAS pools were down, was fixed by 1:30.
- NLT1: ntr
- ASGC: nta
- OSG: ntr
- FNAL: ntr
- NDGF: security incident is still being investigated
- CNAF: ntr
- RAL: ntr
- IN2P3: ntr
- Grid services:
  - ◆ **CERN LHC VOMS** Yesterday Wed 10th August from around 09:00 to 17:30 voms-proxy-init was frequently failing against voms.cern.ch and lcg-voms.cern.ch a lot more often than *normal*. The error would have been **Error: lhcb: Problems in DB communication**. Apparently the normally excessive redundancy in the VOMS service did the



job and most people did actually get proxy AFAIK. The daemons were in a *bad state* which has been diagnosed and reported GGUS:73383 <http://cern.ch/go/78rj>

- ◆ **CERN LHC VOMS** , The upgrade from from gLite 3.2 VOMS to EMI 1 VOMS is now possible. Proposal:
  - ◇ Wednesday 17th August 10:00 -> 12:00 CEST , upgrade of voms.cern.ch
  - ◇ Thursday 18th August 10:00 -> 12:00 CEST , upgrade of lcg-voms.cern.ch
  - ◇ The migration should be completely transparent and rollback is easily possible.
  - ◇ There are no expected changes in behaviour except:
    - The ugly voms-proxy-init miss messages improves from **Suspension reason: ORA-01062**: to something readable.
    - The oracle session cursor leak per voms-proxy-init miss and eventual server crash is expected to be fixed.
  - ◇ This intervention will be confirmed and published after Friday August 12th 15:00 meeting.
- Storage services:
  - ◆ **CERN CASTOR** - would like to start planning the interventions during the next technical stop. Will send more detailed proposals to each experiment concerned (and then do announcements via GOCDB etc)
    - ◇ CASTORCMS: 2011-08-25 (afternoon): 30min downtime (switch from LSF to transfermanager, update to 2.1.11-2)
    - ◇ CASTORATLAS: 2011-08-29: 2h downtime (to be confirmed); database defrag + change database storage HW + update to 2.1.11-2
    - ◇ CASTORLHCB: 2011-08-30: 30min downtime (switch from LSF to transfermanager, update to 2.1.11-2)
    - ◇ CASTORALICE: 2011-08-31: 30min downtime (switch from LSF to transfermanager, update to 2.1.11-2)
    - ◇ **all** CASTOR instances: 2011-08-31: 2h downtime (new HW for central nameserver database)
  - ◆ **CERN CASTOR** - upcoming transparent interventions (will get announced to the experiment) and interventions not directly affecting the LHC experiments (will just go to CERN statusboard)
    - ◇ transparent update to all CASTOR xroot redirectors (service restart, covered by client retry).
    - ◇ transparent change on CASTOR-SRM servers (use per-instance CASTOR nameservers instead of central nameservers)
    - ◇ CASTORPUBLIC: 2011-08-17: 1h downtime (change database storage HW, switch from LSF to transfermanager, update to 2.1.11-2). CERN\_PROD will be declared unavailable in GocDB due to SAM SRM tests, but intervention does not directly affect LHC experiments.

AOB: none

## Friday

Attendance: local (AndreaV, Elena, Jamie, Jarka, Ulrich, Jan, Mike, Eva); remote (Alexander, Giovanni, Catalin, Christian, Marc, Andreas, Gareth, Jhen-Wei, Rob; Daniele, Vladimir).

Experiments round table:

- ATLAS reports -
  - ◆ Physics
    - ◇ New Physics Run since last night
  - ◆ T0/Central Services

- ◊ High load on ATLR DB problem: after discussion it was decided to isolate T0 load on the database from the COOL and PVSS data replication PVSS2COOL data copying. IT DBAs is going today to split ATLR : node 1 will serve the Oracle Streams replication processes of COOL and PVSS data, PVSS2COOL processes, PVSS configuration the archive reading; nodes 2 and 3 will serve T0 jobs (ATLAS\_COOL\_READER\_TZ) (up to now these sessions where on all 3 nodes of the cluster) and ATLAS\_COOL\_READER\_U jobs. [Eva: the database intervention has been completed. Andrea: what has changed to suddenly trigger this problem with COOL and PVSS? Jarka: the trigger rate has increased (LHC delivers more data, and the ATLAS acceptance has also been increased).]
- ◆ T1 sites
  - ◊ INFN-T1: srm down (GGUS:73422,73423). Fixed in 30 min. Thanks.
- ◆ T2 sites
  - ◊ Nothing to report
- CMS reports -
  - ◆ LHC / CMS detector
    - ◊ Fills for physics
  - ◆ CERN / central services
    - ◊ CASTORCMS: SLS monitoring drops in availability, GGUS opened by S.Gowdy (GGUS:73430). Further commented by D.Bonacorsi to point to CASTORCMS\_CMSCAFUSER.
    - ◊ CASTORCMS\_DEFAULT: several short SLS monitoring drops in availability, labelled as "glitches", no tickets opened. The same for T0EXPRESS, CMST3, and others.
    - ◊ [Jan: about CASTORCMS\_DEFAULT, this is a known issue. Note also that this morning some problems with SLS were observed around 11am, which are probably related to the changes in checksum verification performed at 9am. These changes should only affect idle disks, but maybe are not working as expected - please report any issue you should observe. Daniele: thanks, will report it to CMS operations.]
  - ◆ T1 sites:
    - ◊ PIC: since the downtime last Tuesday, the glideIn WMS pilots cannot enter the site via ce07.pic.es, held in the factory's condor queue. PIC admins report they have recently changed their gridmapdir, (Savannah:122802, GGUS:73414)
    - ◊ ASGC: transfer issues to UCSD and Vienna reported yesterday have been understood and fixed, tickets CLOSED (Savannah:122785, Savannah:122786).
  - ◆ T2 sites:
    - ◊ a couple of sites failing bdII visibility tests: T2\_DE\_RWTH (Savannah:122809), T2\_AT\_Vienna (Savannah:122808). Both now OK, both ~14 hrs long, roughly same time window (would this point to bdII more than to sites?)
- ALICE reports -
  - ◆ Nothing to report
- LHCb reports -
  - ◆ Experiment activities:
    - ◊ Data processing and stripping. We still have backlog of jobs at GRIDKA.
  - ◆ New GGUS (or RT) tickets:
    - ◊ T0: 0
    - ◊ T1: 0
    - ◊ T2: 0
  - ◆ Issues at the sites and services
    - ◊ T0
      - The CASTOR glitch reported yesterday (INC:058607) was understood: due

to some incorrectly finished operation we had few files with zero size at CASTOR. We tried infinitely to send these files to CERN and all these transfers failed. After removing these corrupted files we have no any problems.

◇ T1

· ntr

Sites / Services round table:

- NLT1: ntr
- CNAF: ntr
- FNAL: ntr
- NDGF: two points
  - ◆ We believe that the security incident in Oslo had no consequences, but will reinstall all machines anyway. This should be done by next Monday.
  - ◆ There will be an intervention in the main NDGF site on Monday morning at 10am.
- IN2P3: ntr
- KIT: ntr
- RAL: ntr
- ASGC: ntr
- OSG: ntr
  
- Grid services: ntr
- Database services: nta
- Dashboard services: ntr
- Storage services: checksum verification has now been turned for all users (already done yesterday for ATLAS and LHCb).

AOB: none

-- JamieShiers - 18-Jul-2011

---

This topic: LCG > WLCGDailyMeetingsWeek110808

Topic revision: r16 - 2011-08-12 - AndreaValassi



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