

Table of Contents

Week of 130121.....	1
Daily WLCG Operations Call details.....	1
WLCG Availability, Service Incidents, Broadcasts, Operations Web.....	1
General Information.....	1
Monday.....	1
Tuesday.....	3
Wednesday.....	4
Thursday.....	6
Friday.....	7

Week of 130121

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 Availability, Service Incidents, Broadcasts, Operations Web

VO Summaries of Site Usability				SIRs	Broadcasts	Operations Web
ALICE	ATLAS	CMS	LHCb	WLCG Service Incident Reports	Broadcast archive	Operations Web

General Information

General Information			GGUS Information	LHC Machine Information
CERN IT status board	WLCG Baseline Versions	WLCG Blogs	GgusInformation	Sharepoint site - LHC Page 1

Monday

Attendance: local(Simone, Ignacio, Xavi, Belinda, Maarten - Alice, Ian - CMS, Stefan - LHCb, Eddie, Emmanuel, Maria D.); remote(Doug - ATLAS, Wei-Jen - ASGC, Saverio - CNAF, Michael - BNL, Onno - NLT1, Matteo - CNAF, Gonzalo - PIC, Tiju - RAL, Rolf - IN2P3, Pepe - PIC, Miguel - NDGF)

Experiments round table:

- ATLAS reports -
 - ◆ CERN -
 - ◇ CERN CA CRL expired on 19 Jan , renewed at 14:18 (see [GGUS:90599](#))
 - ◇ CERN LFC affected by bad CRL [GGUS:90602](#)
 - ◇ CERN SRM also affected, service erratic until 20-Jan early in morning [GGUS:90605](#)
 - ◇ Panda servers affected by bad CRL (refreshed CRL by hand) - significant job loss (~30K)
 - ◆ RAL-LG2 affected by bad CRL [GGUS:90596](#) and [GGUS:90589](#)
 - ◆ TRIUMF affected by bad CRL [GGUS:90594](#)
 - ◆ After CRL refresh all services back to normal
- CMS reports -
 - ◆ LHC / CMS
 - ◇ Beam over weekend
 - ◆ CERN / central services and T0
 - ◇ Lost the CRL list for the CERN CA on Saturday afternoon. Alarm tickets issued because all certs then listed as expired. Perhaps a more accurate error message could

- ◊ be created?
- ◊ Parallel Tier-0 seems to be working. Working on Validation with the eye toward shutting off the old.
- ◆ Tier-1:
 - ◊ NTR
- ◆ Tier-2:
 - ◊ NTR
- ALICE reports [↗](#) -
 - ◆ CERN:
 - ◊ CERN CA CRL expired on Sat (GGUS:90599 [↗](#))
 - ◊ ca-proxy.cern.ch kept serving expired CRL for CERN CA (GGUS:90600 [↗](#))
 - ◆ KIT:
 - ◊ GGUS web server kept expired CRL for CERN CA (GGUS:90601 [↗](#))
- LHCb reports -
 - ◆ Mostly Simulation jobs on all Tier levels
 - ◆ 2011 data reprocessing to be started today/tomorrow
 - ◆ T0:
 - ◊ VOMS failed on the week-end, many grid services affected, (no GGUS b/c need certificate for ticket submission), many thanks for prompt resolution by IT/PES after calling operator who sent SMS to support team.
 - ◊ After change of the GRIDKA SRM endpoint also monitoring (SLS,SUM) have been updated accordingly
 - ◆ T1:
 - ◊ GRIDKA: transfers to be watched, as there is very little data currently transferred also no errors. pA data will be replicated to GRIDKA this week which will be a good test for FTS.

Sites / Services round table:

- CERN CA CRL out-of-date:
 - ◆ Emmanuel: last tuesday, a new version of the CERN CA website was deployed. This broke the location of the web site for the CRL file. The internal CERN IT monitoring did not detect it, nor did the one of EUGridPMA (the reason of this still being investigated). The first alert come on saturday and the CRL file was upgraded before expiration, but the propagation in caches on the Grid took time and therefore the CRL file expired in many Grid locations. The CERN CA experts will improve the monitoring to have a check every hour. Also, it is probably time to review the mechanism of Grid caching every 6 hours. This can also be a security issue in case of compromised credentials (Maarten will discuss with security people). A SIR will be provided.
 - ◆ Maarten, concerning the CMS report: the issue of error messages has been raised at the time of the WLCG Ops TEG. It will be further discussed with the developers of low level components. In addition, a discussion about expired CRLs and the possibility to have only warning rather than fatal failures was started in the Security TEG.
 - ◆ Maarten: the CERN proxy serving CRLs was not updated after many hours, not clear if due to the caching of 6 hours mentioned above or to some other problem (it will be investigated further and a statement will be made in the SIR). The cache was refreshed by hand. The GGUS web server did not upgrade even in the evening (so there was for sure a problem there, still to be understood). Maria D reminded that the authorized alarmers can login with username and password in GGUS and submit alarm in case the certificate is not accepted (as it happened to LHCb during the incident, CMS and ATLAS used a non-CERN certificate, Maarten logged in with user name and password).

- ASGC: CASTOR Database crashed again, now OK.
- CERN:
 - ◆ Belinda: EOSATLAS being updated. CMS upgrade will be planned soon
 - ◆ Ignacio: the LFC upgrade to the EMI release needs to be planned. Alessandro: discussed with ATLAS, the upgrade can go ahead. Ignacio: the intervention should be transparent, will start tomorrow. LHCb will go after.
- MariaD/GGUS: Reminder: The January release will take place this Wednesday on 2013/01/23. There will be the usual series of test ALARMS.

AOB:

Tuesday

Attendance: local(Simone,Xavi, Guido - ATLAS, Ignacio, Eddie, Belinda, Maria D.); remote(Michael - BNL, Saverio - CNAF, Xavier - KIT, Miguel - NDGF, Ronald - NL-T1, Wei-Jen - ASGC, Joel - LHCb, Lisa - FNAL, Tiju - RAL, Rob - OSG, Rolf - IN2P3, Ian - CMS)

Experiments round table:

- ATLAS reports -
 - ◆ CERN
 - ◇ EOS update on monday afternoon (both eos and xrootd software were updated) - no problem reported
 - ◆ Tier1s: NTR
 - ◆ Tier2s
 - ◇ RO-16-UAIC: jobs failing with ""lost heartbeat" (GGUS:90665 [↗](#)), could be due to a site reconfig. Situation seems better
 - ◇ MWT2_UC: errors in transfers (GGUS:90620 [↗](#)), no comments from the site after 24h, but problems disappeared
 - ◇ INFN-MILANO: errors in transfers (GGUS:90597 [↗](#)), problem with log rotation in the storm frontend, the sys admins restarted it
- CMS reports -
 - ◆ LHC / CMS
 - ◇ Running
 - ◆ CERN / central services and T0
 - ◇ Castor slowdown issue over night. Team ticket submitted
 - ◇ Operations team would like an explanation of the acronyms in the following plot. https://lemon.cern.ch/lemon-web/index.php?target=get_rrd_graph_advanced&graph_type=ma
 - ◇ (from CASTOR team, new monitoring in place, please use: <http://c2adm01.cern.ch/metric/RecallBacklogSize-CMS> [↗](#))
 - ◆ Tier-1:
 - ◇ <https://savannah.cern.ch/support/index.php?135288> [↗](#) Failure to submit to FTS at KIT. (I tried to retroactively bridge the ticket to GGUS, which seems not to have worked)
 - ◆ Tier-2:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ Mostly Simulation jobs on all Tier levels

- ◆ 2011 data reprocessing to be started today/tomorrow
- ◆ T0: CERN : LFC upgrade on Thursday
- ◆ T1:

Sites / Services round table:

- OSG Operations Maintenance window that is running today, See:
<http://osggoc.blogspot.com/2013/01/goc-service-update-tuesday-january-22nd.html>
- PIC: <https://savannah.cern.ch/support/?135266> -> The CVMFS seems not to catch a configuration file from PIC CMS SITECONF, and this has translated into the failure of the cms-basic SAM test. The changed was made last Friday, and we spent all the weekend in 'warning' state. The test became 'critical' yesterday. This is not a site problem, but rather the sync. of CVS config files to CVMFS seems to be broken. CMS has been notified and we are waiting for it to be fixed.
- GGUS:
 - ◆ Slides with GGUS activity and real ALARM drills are attached to this pages. File ggus-tickets.xls is up-to-date and attached to page WLCGOperationsMeetings. It contains 4 years of WLCG tickets. We shall chop this file when the LHC run stops next month.
 - ◆ ALARM tests tomorrow due to the release. Supporters, site managers, please **do not verify** the tickets before we are sure that their diary contains entries by:
 1. the operators (when applicable)
 2. an experiment member in e-group -operator-alarm@cern.ch (when applicable)
 3. a service manager of the (supposedly) affected service.
- CERN:
 - ◆ CA incident on Saturday: Please find the SIR in <https://twiki.cern.ch/twiki/bin/view/PESgroup/IncidentCA2013>
 - ◆ LFC upgrade to the EMI: The DB schema update has been announced in GOCDB for Wednesday 23 Jan on ATLAS and for Thursday 24 Jan on LHCb and shared LFC instances. Please note that we will be adding LFC front-end nodes with the new software shortly after the schema update and if things look OK, reinstall the older nodes progressively afterwards. All this should be transparent to LFC users.
 - ◆ Xavi: CASTOR DB backend loaded yesterday, investigating together with DB people - not clear picture yet. There is a backlog in castor of 50TB to be recalled from a couple of CMS users (ceballos, paus) in CMS default pool. There is a new CASTOR monitoring which is linked below.
 - ◆ Belinda: Increased quota for EOSLHCb to 800TB.
 - ◆ Eddie: ATLASDDM 1.0 dashboard has been degraded since 11:00 PM yesterday due to DB issues (exceptional read load). Back in production since 3:00PM. ATLAS DDM dashboard 2.0 (new version) not affected.

AOB:

Wednesday

Attendance: local(Simone, Xavi E., Guido - ATLAS, Luca C., Ignacio, Eddie, Maria, D. Belinda, Maarten); remote(Pepe - PIC, Stefano - LHCb, Saverio - CNAF, Michael - BNL, Miguel - NDGF, Matteo - CNAF, Alexander - NL-T1, Wei-Jen - ASGC, Rolf - IN2P3, Tiju - RAL, Lisa - FNAL, Ron - NL-T1, Pavel - KIT, Rob - OSG)

Experiments round table:

- ATLAS reports -
 - ◆ Services

- ◊ upgrade of the schema in ATLAS LFC
(https://goc.egi.eu/portal/index.php?Page_Type=View_Object&object_id=120133&grid_id=0)
- ◊ transfers of files >5GB still failing (now at DESY-HH); ticket https://ggus.eu/tech/ticket_show.php?ticket=89998 not totally clear: can the patched version be tested/deployed?
- ◆ Tier0/Tier1s: NTR
- ◆ Tier2s
 - ◊ NERSC showing lots of errors in receiving data (GGUS:90619)
- CMS reports -
- ALICE reports -
 - ◆ NTR
- LHCb reports -
 - ◆ 2011 data reprocessing started yesterday and extended today morning. First jobs already went through all the steps
 - ◆ Share of CPU resources changed:
 - ◊ 80% of reprocessing performed at T1s and the rest at CERN
 - ◊ Processing of current pA data taking shared equally between GRIDKA and CERN
 - ◆ T0:
 - ◊ LFC upgrade to EMI2 tomorrow
 - ◊ GGUS:90752 - Jobs failing because the script used to setup the environment are failing. It seems a timeout problem connected with CVMFS. We tried to login on the affected WN's but we did not manage.
 - ◆ T1:

Sites / Services round table:

- OSG: maintenance mentioned yesterday went smooth.
- PIC: the CVMFS problem at PIC mentioned yesterday due to a problem in a CMS configuration file was fixed by CMS people.
- CERN
 - ◆ DB: about the issue with CASTOR, it was the side effect of a bug (to be understood) introduced by a change on monday which introduced a parameter to better handle data corruption. The parameter is used in many other databases at CERN but the issues were observed only for CASTOR DB. The parameter was reverted.
 - ◆ DB schema upgrade to LFC was OK (no issues observed). But some GSS errors have been seen in the new LFC frontends with EMI2, so upgrade of frontends to new EMI2 release is stopped for moment, waiting for input from LFC developers.
- GGUS:
 - ◆ About IanF's comment yesterday on Savannah:135288, according to the https://wiki.egi.eu/wiki/GGUS:Savannah_Bridge_FAQ all fields are mapped correctly and GGUS:90688 is created. The savannah original submission text is present in the GGUS 'Detailed description'. In order to have a better format, the flag 'Use GGUS:' should have been set from the beginning, according to the original agreement with CMS.
 - ◆ Experiments or sites with GGUS tickets **not** receiving adequate support, please send them to MariaD a.s.a.p. for presentation at tomorrow's Ops Coord Meeting <https://indico.cern.ch/conferenceDisplay.py?confId=231008>
 - ◆ Please note the <https://ggus.eu/pages/didyouknow.php#2013-01-23> news about new functionality on TEAM / ALARM members' list for one's own VO.

- ◆ Today's release test ALARMS GGUS:90724 [↗](#) - GGUS:90736 [↗](#) (all intermediate numbers) 13 tickets by 10am UTC went well.

AOB:

- SIR CASTOR incident

Thursday

Attendance: local(Simone, Ueda, Luc Poggioli - ATLAS, Ignacio, Belinda, Eddie, Maarten - Alice); remote(Wei-Jen - ASGC, Saverio - CNAF, Michael - BNL, Stefano - LHCb, Matteo - CNAF, Ronald - NL-T1, Jeremy - GridPP, John - RAL, Kyle - OSG, Marian - KIT, Miguel - NDGF, Rolf - IN2P3, Ian - CMS, Lisa - FNAL, Pepe - PIC)

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central services
 - ◇ Reading problems from castoratlas. GGUS:90759 [↗](#) (Alarm initially). Controller pb on a given server. List of corrupted files to be provided
 - ◇ Transfers of files >5GB still failing (now at DESY-HH); GGUS:89998 [↗](#). WLCG broadcast to deploy the patch or wait for EMI release?
 - ◆ Tier1s
 - ◇ RAL Backlog of functional test transfers
 - ◆ Tier2s
 - ◇ NTR
- CMS reports -
 - ◆ LHC / CMS
 - ◇ Running
 - ◆ CERN / central services and T0
 - ◇ Validation of the new T0 going well. Switching from the old to the new T0 machinery T0 will reduce the load to CASTOR by 25%.
 - ◆ Tier-1:
 - ◇ Consistency Check requests were sent to all Tier-1s
 - ◆ Tier-2:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ KIT: some Xrootd servers could not be accessed by jobs, due to network or Xrootd SW issues; experts looked into it and restarted a few daemons.
- LHCb reports -
 - ◆ Very intense activity: reprocessing, prompt-processing, MC and user jobs. We are hitting our limits in running jobs.
 - ◆ T0:
 - ◇ LFC upgrade to EMI2 ongoing
 - ◇ Problem in transferring job output from the pit to EOS. We are able to ping the EOS system but transfer fails.
 - ◇ INC:226288 [↗](#) - Started tests for data replication from Castor to EOS. The throughput is not as good as expected.
 - ◇ Stratum 1 of CVMFS all report degraded status in SLS.
 - ◆ T1:

◇ RAL: some jobs failing because not able to setup the environment (timeout of the script). The problem is not affecting some particular WNs. Not a major issue. Already in touch with the contact person. They are investigating.

Sites / Services round table:

- KIT: general routing problem for LHCOPN at KIT. Fully recovered only this morning. Could explain also Alice problems reported about xrootd. Announced at RISK for all SEs from 8 to 10 UTC on Tuesday for a restart of SW layer to tape connection.
- CERN: for the LFC upgrade, the schema upgrade has gone well but the deployment of frontends stopped because of a GSS error. The problem can be solved by pointing to an Oracle10 client on AFS, which is not ideal. During the meeting we agreed there is no rush upgrading to the EMI release, so the intervention will be stopped until the LFC developers can provide a patch.

AOB:

Friday

Attendance: local(Simone, Luc Poggioli - ATLAS, Belinda, Luca M., Ignacio, Stefano - LHCb, Eddie, Eva, Maarten); remote(Jeremy - GridPP, Michael - BNL, Xavier - KIT, Lisa - FNAL, Wei-Jen - ASGC, Gareth - RAL, Onno - NL-T1, Rob - OSG, Rolf - IN2P3, Ian - CMS, Matteo - CNAF, Ulf - NDGF)

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central services
 - ◇ Reading problems from castoratlas. GGUS:90759 [↗](#). (Controller pb on a given server). List of lost file (7) available.
 - ◇ AFS frozen. GGUS:90796 [↗](#) & <http://itssb.web.cern.ch/service-incident/problem-afs154/24-01-2013> [↗](#). HW pb. Solved very fast.
 - ◆ Tier1s
 - ◇ RAL Backlog of functional test transfers due to FTS3 misconfig. SAV:135151 [↗](#). OK now.
 - ◇ PIC: GGUS:90779 [↗](#) Install missing since switch to CVMFS. Environment variables corrected. Fixed.
 - ◆ Tier2s
 - ◇ NTR
- CMS reports -
 - ◆ LHC / CMS
 - ◇ Running
 - ◆ CERN / central services and T0
 - ◇ Team ticket issued for Castor, when a slow down from P5 was observed
 - ◆ Tier-1:
 - ◇ Hammer Cloud failure reported at ASGC
 - ◆ Tier-2:
 - ◇ NTR
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -

- ◆ Activity is going as yesterday Very intense activity: reprocessing, prompt-processin, MC and user jobs. We are hitting our limits in running jobs.
- ◆ T0:
 - ◇ LFC upgrade to EMI2 -> some issue related to Oracle version. Investigating
 - ◇ Transferring job output from the pit to EOS -> still investigating.
 - ◇ INC:226288 - After some tuning the throughput increased. Still to evaluate if it is enough.
 - ◇ Comment from Luca M. (CERN storage): for EOS, the only test done was a traceroute, it would be nice to try some real transfer (xrdcp and gridftp). For the FTS the issue was with the low number of concurrent configured transfers. There was also a problem with the new CRL for one EOS disk server since it was a new machines.
- ◆ T1:

Sites / Services round table:

- ASGC: This morning there were some jobs failed on our CEs. That's because we forgot to turn off yum-autoupdate daemon so they downloaded some wrong packages. The problem has been fixed and jobs work fine now.
- KIT: Alice is transferring a lot of data from WN to world. The KIT firewall is suffering. Maarten will pass the message. KIT is being used for the prompt reconstruction.
 - ◆ ATLAS for KIT: there is an ongoing issue for FZK staging from TAPE from last october. There is a quite old ticket still open, would be good to have a comment. Answer from KIT: the AT RISK scheduled next week is for an intervention to improve the situation.
- RAL: AT RISK on monday for power input work.
- CERN: For the LFC upgrade, there should be a build with Oracle11 by today. KIT will need this as well.

AOB:

This topic: LCG > WLCGDailyMeetingsWeek130121

Topic revision: r39 - 2013-01-25 - MaartenLitmaath



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