

Table of Contents

| | |
|---|----------|
| Week of 130211..... | 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 130211

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 (AndreaV, Alexandre, Rajan, Xavi, Peter, Eva, Maarten, Ulrich); remote (Wei-Jen/ASGC, Saverio/CNAF, Ronald/NLT1, Christian/NDGF, Tiju/RAL; Torre/ATLAS).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central services
 - ◇ AFS sls not available. [INC:236796](#). Back to normal very fast
 - ◇ LSF: many EOWYN heartbeat + Recon issues.
 - ◇ LSF does not respond. [GGUS:91329](#) (Alarm). Restarted fast, but scheduling/dispatching problems. Fixed. [Ulrich: this is essentially the same issue reported by CMS.]
 - ◆ Tier1s
 - ◇ FZK-LCG2 failing transfers. [GGUS:91316](#) (Alarm). Routing problem (affected also CMS). Solved very fast
 - ◇ CNAF, FZK, NIKHEF still out of T0 export (DATADISKS almost full)
 - ◆ Calibration-Tier2s
 - ◇ IFIC-LCG2 SRM issue. [GGUS:91327](#). A priori solved. [Torre: having troubles with GGUS at the moment so cannot check. There seems to be some issue with the DoE grid. Maarten: can you access GOCDB at [goc.egi.eu](#)? Torre: will follow up after the meeting.]
- CMS reports -
 - ◆ LHC / CMS

- ◊ CMS ready and waiting to acquire low energy pp data
- ◆ CERN / central services and T0
 - ◊ ALARM ticket (GGUS:91325 [↗](#)) opened to CERN-PROD on 2013-02-09 19:36 UTC, because the main CMS Tier-0 node had been down for 7h (originally TEAM but then escalated to ALARM 20 minutes later). Rebooted 2013-02-09 19:56 UTC. Note that CMS didn't see any lemon alarm because the machine (vocms15) was in maintenance mode, which originally was a mistake on the CMS side. It is not clear tho why the machine went completely offline. [Maarten: is the node in production now? Peter: yes it is in production.]
 - ◊ ALARM ticket (GGUS:91328 [↗](#)) opened to CERN-PROD on 2013-02-10 08:46 UTC, because no LSF submission were possible since 2013-02-10 06:37 UTC (originally TEAM but then escalated to ALARM 1 hour later). After the intervention by CERN/IT on 2013-02-10 10:03 UTC, the situation improved. Note that the Hammercloud tests against the CMS Tier-0 went back to green only 3-4 hours after the intervention, and back to red again during the night (Feb 10-11). The CMS shift crew opened another ticket (GGUS:91335 [↗](#)) which could have been avoided, since it had been indicated on the CERN/IT SSB that after the master batch daemon crashed and had to be manually restarted, a reconfiguration was necessary to resolve a subsequent dispatch issue, which probably caused the CMS HC issue. We apologize for that subsequent ticket. Now everything is back to green, thank you. [Ulrich: LSF went down due to the crash of a process, everything was fixed but it took some time. A SIR will be prepared. Peter: is this related in any way to the maximum nuber of users in LSF? Ulrich: no this is not related. Peter: did this only affect CMS? Ulrich: no this was a general issue.]
- ◆ Tier-1:
 - ◊ T1_DE_KIT had a SAM SUM SRM failure 2013-02-08 21:41 UTC (GGUS:91317 [↗](#)) : identified by the network department at KIT as a "very obscure routing effect, which is triggered with a weird time lag". Fixed on 2013-02-09 11:41 UTC.
- ◆ Tier-2:
 - ◊ NTR
- ◆ AOB
 - ◊ NTR
- ALICE reports [↗](#) -
 - ◆ CERN: job submission to the CREAM CEs has often been very slow in the last few days, leading to a large shortfall in the use of CERN resources by ALICE, as the submission could not keep up with the rate of jobs finishing. As of ~13:00 today things look normal again, but is the problem understood? [Maarten: this is not related to the LSF issue discussed for CMS and ATLAS. Ulrich: will analyse this after the meeting, please open a ticket.]
 - ◆ [Tiju/RAL: last week we got "no space left on device" issues for ALICE. Maarten: will forward the info to the ALICE colleagues in charge of storage. Will also investigate if Monalisa sent reports to the ALICE experts, this is an issue that can only be solved by ALICE because they manage the space, the site cannot do anything about it even if they are notified.]
- LHCb reports -
 - ◆ Ongoing activity as before: reprocessing, prompt-processing, MC and user jobs.
 - ◆ T0:
 - ◊ NTR
 - ◆ T1:
 - ◊ IN2P3: NAGIOS problem still being investigated (GGUS:91126 [↗](#)). Also, low level problem with access to data (input data resolution) - under investigation by IN2P3 contact.
 - ◊ RAL: Continuing problems with batch system (GGUS:91251 [↗](#))
 - ◊ FZK : Problem with FTS transfers solved over the weekend (GGUS:91315 [↗](#)).

Sites / Services round table:

- Wei-Jen/ASGC: ntr
- Saverio/CNAF: ntr
- Ronald/NLT1: ntr
- Christian/NDGF: downtime this morning to move pools to OPN, all went OK
- Tiju/RAL: on Wednesday we will permanently turn off AFS client access from our worker nodes
- Michael/BNL [via email]: ntr
- Pepe/PIC [via email]: ntr
- Kyle/OSG [via email]: ntr

- Ulrich/Grid: pre-announcement, plan to continue CE upgrade to EMI2 soon
- Eva/Databases: ntr
- Alexandre/Dashboard: ntr
- Xavi/Storage: yesterday a router glitch in the CC briefly affected around 400 of our nodes

AOB: none

Tuesday

Attendance: local(AndreaS, Peter/CMS, Raja/LHCb, Alexandre, Xavi, MariaD, Ulrich);remote(Onno/NL-T1, Torre/ATLAS, Tiju/RAL, Rolf/IN2P3, Matteo/CNAF, Wei-Jen/ASGC, Rob/OSG, Christian/NDGF).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central services
 - ◇ NTR
 - ◆ Tier1s
 - ◇ CNAF, FZK, NIKHEF still out of T0 export (DATADISKS almost full)

- CMS reports -
 - ◆ LHC / CMS
 - ◇ 3 low energy LHC pp fills since last night, successfully processed by CMS (nearly 1 pb⁻¹ delivered Lumi)
 - ◆ CERN / central services and T0
 - ◇ (yet another) ALARM ticket (GGUS:91380) opened to CERN-PROD on 2013-02-11 22:54 UTC. Due to a critical CASTOR file reading degradation (T0Export Pool) while the pp run was about to start (TEAM ticket escalated to ALARM). After prompt response by the CASTOR team, it turned out the issue was due to 1 dataset that was heavily accessed by a CMS Tier-0 Replay workflow. The reason was that the dataset in question came from a high rate run from 2012, while the Replay workflow was using a standard job splitting configuration, hence triggering a very high number of jobs on that dataset. Apparently the issue is still affecting the CMS Tier-0 at the time of writing this report, but we have enough experts on the case. [Xavi mentions that there was a second problem due to one of the head nodes having an expired certificate for an xrootd redirector.]
 - ◆ Dashboard (Site Status Board) : CMS opened a ticket (Savannah:135851) for a couple site downtimes not being reported neither in the CMS SSB nor in the CMS Site Downtime GCal (same data source). The Dashboard team responded this is the consequence of known bugs (BUG:99497 and BUG:100157) for which fixes were found and will be deployed on the production CMS SSB next Tue/Wed. Until then, CMS is kindly asking the Dashboard team, if not too painful for them, to insert manually all known scheduled Site Downtimes in the SSB, Thank You !

- ◆ Tier-1:
 - ◇ NTR
- ◆ Tier-2:
 - ◇ NTR
- ◆ AOB
 - ◇ Peter Kreuzer extended his CRC-term until the end of the on-going run (Thu Feb 14, 8AM)
 - ◇ From this Thursday the CMS report will often be given remotely

Rolf reports that since at least January the GGUS tickets generated from Savannah do not have anymore the VO field. Maria is not aware of the problem and will investigate.

- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ Ongoing activity as before: reprocessing, prompt-processing, MC and user jobs.
 - ◆ T0: NTR
 - ◆ T1:
 - ◇ IN2P3: NAGIOS problem still ongoing (GGUS:91126 [↗](#)). No idea who to follow up with. [Rolf adds that from the IN2P3 there is absolutely nothing strange with the Nagios jobs, they appear to run and complete normally. The problem must be in the test output. Andrea will give a look, as in the past he saw a similar problem for CMS.]
 - ◇ RAL: Problems with batch system seem to be resolved (GGUS:91251 [↗](#))

Sites / Services round table:

- ASGC: this morning FTS went down and was restarted.
- BNL: all Tier-1 services are fully operational and being monitored by the Tier-1 personnel. The lab was closed today for the facility operations team to complete the snow removal.
- CNAF: ntr
- IN2P3: ntr
- NDGF: ntr
- NL-T1: on February 25-26 SARA will have a major network intervention and in addition the WNs will be migrated to EMI-2 and SL6. [Andrea: as far as I know ATLAS is not yet OK with Tier-1 sites moving to SL6 as some workflows are not yet fully validated. Onno will check with ATLAS.]
- RAL: ntr
- OSG: just to inform that BNL is closed today, this explains why they are not connected. [See BNL report.]
- CERN batch and grid services: ntr
- CERN storage services: ntr
- Dashboards: ntr
- GGUS: File ggus-tickets.xls is up-to-date on page WLCGOperationsMeetings with GGUS totals per experiment. We had 7 real ALARMS since the last MB.

AOB:

Wednesday

Attendance: local (AndreaV, Xavi, LucaC, Rajan, Eddie, Ulrich); remote (Michael/BNL, Alexandre/NLT1, Saverio/CNAF, Dimitrios/RAL, Pavel/KIT, Lisa/FNAL, Christian/NDGF, Wei-Jen/ASGC, Rob/OSG, Rolf/IN2P3; Peter/CMS, Torre/ATLAS; MariaD/GGUS).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central services
 - ◇ NTR
 - ◆ Tier1s
 - ◇ NTR
- CMS reports -
 - ◆ LHC / CMS
 - ◇ 2 long low energy LHC pp fills since yesterday evening (>3 pb⁻¹ delivered Lumi)
 - ◆ CERN / central services and T0
 - ◇ Regarding the ticket opened on Feb 11 22:54 for CASTOR file reading degradation on the T0Export pool (GGUS:91380): the situation improved as CMS significantly reduced the load (number of running jobs). Then the number was increased again, but the farm wasn't really stressed significantly either yet. So the immediate problem is gone, and therefore the ticket can be closed, but CMS needs closer analysis on what happened for the longer term future.
 - ◆ Tier-1:
 - ◇ NTR
 - ◆ Tier-2:
 - ◇ NTR
 - ◆ AOB
 - ◇ Peter Kreuzer extended his CRC-term until the end of the on-going run (Thu Feb 14, 8AM)
- ALICE reports -
 - ◆ NTR
- LHCb reports -
 - ◆ Ongoing activity as before: reprocessing, prompt-processing, MC and user jobs.
 - ◆ T0: NTR
 - ◆ T1:
 - ◇ IN2P3: NAGIOS problem still ongoing (GGUS:91126).
 - ◇ RAL: Problems with batch system came back (GGUS:91251)
 - ◇ GridKa : Possible issue with srm / SE / network (GGUS:91474). Jobs failing to resolve input data multiple times at GridKa. Jobs at JINR waiting for a long time for data from GridKa, before being killed by the batch system there.

Sites / Services round table:

- Michael/BNL: ntr
- Alexandre/NLT1: one fileserver crashed this morning and was restarted, all OK now
- Saverio/CNAF: ntr
- Dimitrios/RAL: as planned, we permanently disabled AFS on our batch nodes today
- Pavel/KIT: problem on many WNs caused by memory leaks in ALICE user jobs, the user has been contacted and this is being fixed
- Lisa/FNAL: ntr
- Christian/NDGF: ntr
- Wei-Jen/ASGC: ntr
- Rob/OSG: ntr
- Rolf/IN2P3: ntr
- Pepe/PIC [via email]: ntr

- LucaC/Databases: ntr
- Eddie/Dashboard: ntr
- Xavi/Storage: ntr
- Ulrich/Grid: ntr
- MariaD/GGUS: ntr

AOB: none

Thursday

Attendance: local (AndreaV, Raja, Alexandre, MariaD, Ulrich); remote (Matteo/CNAF, John/RAL, Wei-Jen/ASGC, Marian/KIT, Ronald/NLT1, Lisa/FNAL, Pepe/PIC, Michael/BNL, Jeremy/Gridpp, Rolf/IN2P3, Rob/OSG, Christian/NDGF; Torre/ATLAS).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central services
 - ◇ T0 LSF dispatching: ALARM ticket 21:03 last night. Pending jobs accumulating. When an expert looked 90min later, pending jobs had cleared. Probably too low a threshold for alarm ticket, especially hours before run ends. GGUS:91501 [↗](#) [Ulrich: the system was already recovering by itself when we started investigating yesterday evening]
 - ◇ Weird behavior for lcg-cp with --srm-timeout option, now seen at RAL also. GGUS:91223 [↗](#) [Raja: we have seen something similar in LHCb, this seems to happen only on SLC6]
 - ◆ Tier1s
 - ◇ Taiwan-LCG2: transfer errors to DATATAPE, ticketed 8:05 this morning. Castor stager DB deadlock, site is working to fix it. GGUS:91505 [↗](#) [Wei-Jen: this is fixed now]
- CMS reports -
 - ◆ LHC / CMS
 - ◇ Run Over
 - ◆ Tier-1:
 - ◇ NTR
 - ◆ Tier-2:
 - ◇ NTR
 - ◆ AOB
 - ◇ CMS CRC has a conflict today for the daily meeting
- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ Ongoing activity as before: reprocessing, some prompt-processing, MC and user jobs.
 - ◆ T0: NTR
 - ◆ T1:
 - ◇ IN2P3: NAGIOS problem still ongoing (GGUS:91126 [↗](#)). [Rolf: someone from LHCb should reassign the ticket. Raja: not sure we can do that either. MariaD: just did it, you could have done it too, wil show Raja how to do it after the meeting.]
 - ◇ RAL: Job timeouts trying to set up environment on the worker node (internal ticket). Continuing problems with batch system (GGUS:91251 [↗](#)).

◇ GridKa : Continuing issue with srm / SE / network (GGUS:91474[↗](#)). Jobs failing to resolve input data multiple times at GridKa. Jobs at JINR waiting for a long time for data from GridKa, before being killed by the batch system there. One strange DNS problem fixed yesterday.

Sites / Services round table:

- Matteo/CNAF: ntr
- John/RAL: ntr
- Wei-Jen/ASGC: nta
- Marian/KIT: ntr
- Ronald/NLT1: ntr
- Lisa/FNAL: ntr
- Pepe/PIC: ntr
- Michael/BNL: ntr
- Jeremy/Gridpp: ntr
- Rolf/IN2P3: nta
- Rob/OSG: we announced at OSG the end of life for the Pacman managing system, we will only use rpm from now on; this should not affect anyone outside OSG
- Christian/NDGF: ntr

- Ulrich/Grid: ntr
- Alexandre/Dashboard: ntr
- MariaD/GGUS: investigating an issue reported by Rolf about CMS tickets originating from savannah and not containing the value CMS in GGUS, this may be due to the December GGUS release

AOB: none

Friday

Attendance: local(AndreaS, Xavi, Alexandre, Przemek, Raja, Ulrich);remote(Torre/ATLAS, Ian/CMS, Michael/BNL, Saverio/CNAF, Rolf/IN2P3, Xavier/KIT, Boris/NDGF, Onno/NL-T1, Pepe/PIC, John/RAL, Jeremy/GridPP).

Experiments round table:

- ATLAS reports -
 - ◆ T0/Central services
 - ◇ Weird behavior for lcg-cp with --srm-timeout option, first seen Taiwan, now seen at RAL also. On SL6, as reported by LHCb. Developer is looking at it. GGUS:91223[↗](#)
 - ◆ Tier1s
 - ◇ BNL: failed transfers from several sources to BNL-OSG2_DATADISK, ticketed 8:36 this morning. GGUS:91548[↗](#) [Michael: we found two overloaded pools and we are now rebalancing the load to prevent transfers from failing. At any rate the fraction of failed transfers was fairly small.]

- CMS reports -
 - ◆ LHC / CMS
 - ◇ Run Over
 - ◆ Tier-1:
 - ◇ NTR
 - ◆ Tier-2:
 - ◇ NTR

- ALICE reports [↗](#) -
 - ◆ NTR
- LHCb reports -
 - ◆ T0: NTR
 - ◆ T1:
 - ◇ IN2P3: Problem with SE at IN2P3 (GGUS:91557 [↗](#))
 - ◇ RAL: Job timeouts trying to set up environment on the worker node (internal ticket).
 - ◇ GridKa : Issue with srm / SE / network resolved (GGUS:91474 [↗](#)) - thanks!
 - ◆ Other : NAGIOS problem still ongoing at IN2P3 (GGUS:91126 [↗](#)). [Andrea reported that before the meeting he saw that the problem seemed to have disappeared in the last 24 hours, but Stefan (contacted after the meeting) thinks this is just by chance.]

Sites / Services round table:

- ASGC: ntr
- BNL: ntr
- CNAF: ntr
- FNAL: ntr
- IN2P3: ntr
- KIT: ntr
- NDGF: ntr
- NL-T1: ntr
- PIC: we are migrating data to the new disks that have been recently installed. As last year we overpledged disk, there will be a decrease in the amount of disk resources for ATLAS (200 TB) and for LHCb (290 TB), while CMS will have an increase. The experiment contacts at PIC are in touch with the experiments.
- RAL: ntr
- GridPP: Some UK T2 sites had problems with multithreaded ATLAS jobs. The user who submitted them was contacted and his jobs were cancelled. Checking if there are other users sending similar jobs. This information may be useful to other sites if they see WN crashes and need a hint to a potential cause, but also for ATLAS to be aware that proof-lite jobs can cause problems.
- CERN batch and grid services: ntr
- CERN storage services: ntr
- Dashboards: ntr
- Databases: ntr

AOB:

- The Alcatel Teamwork system had rather severe problems (INC:240661 [↗](#)). Lisa could not connect because of that.

This topic: LCG > WLCGDailyMeetingsWeek130211

Topic revision: r16 - 2013-02-15 - AndreaSciaba



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