

# Table of Contents

|  |          |
|--|----------|
| <b>Week of 120611.....</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.....   | 2        |
| Wednesday.....   | 4        |
| Thursday.....  | 5        |
| Friday.....  | 7        |

# Week of 120611

## 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 (Andrea, Eric, Maarten, Vladimir, Massimo, David); remote (Rolf/IN2P3, Burt/FNAL, Kyle/OSG, Onno/NLT1, Tiju/RAL, Dimitri/KIT, Zeeshan/NDGF; Ian/CMS, Simone/ATLAS).

Experiments round table:

- ATLAS reports -
  - ◆ T0
    - ◇ On Friday at approx 18:30 EOS started failing any operation. Alarm ticket was submitted to CERN, but EOS people were already looking at the problem (should have checked the IT Status Board). Problem fixed right after.
  - ◆ T1
    - ◇ One disk server in RAL was unavailable starting from friday. The data become available again on saturday. Thanks a lot to the RAL people, those data were urgently needed and would have taken some time to re-generate them.
  - ◆ [Simone: problem with GGUS is still being investigated.]
  - ◆ [Zeeshan/NDGF: downtime at PDC today for ATLAS.]
- CMS reports -
  - ◆ LHC machine / CMS detector
    - ◇ Productive weekend for data taking

- ◆ CERN / central services and T0
  - ◇ We ran short of EOS space at CERN and lead to some number of failures of stage out on repacking, express and reco. This seems to be related to disk servers attempting to accept a file and then filling up and failing. We have freed up space and will adjust the threshold we use to decide the system is full. [Massimo: as we speak, the EOS space for CMS is being increased.]
- ◆ Tier-1/2:
  - ◇ We had a problem with the Site Availability tests over the weekend. We were unable to determine why some sites were failing the SUM CE tests, because there were no reports. Reported and resolved by the dashboard team.
  - ◇ A variety of responses about glxexec tickets last week. We may need some help with the deployment
- ALICE reports [↗](#) -
  - ◆ NTR
- LHCb reports -
  - ◆ Users analysis and prompt reconstruction and stripping at T1s ongoing
  - ◆ MC production at Tiers2
  - ◆ T1:
    - ◇ IN2P3 : Short scheduled Downtime for GE upgrade
  - ◆ [Vladimir: updated wildcard information, but this information is now waiting for validation by USCT. Maarten: please open a GGUS ticket.]

Sites / Services round table:

- Rolf/IN2P3: scheduled downtime at risk for batch, the update went well, all ok now
- Burt/FNAL: ntr
- Kyle/OSG: ntr
- Onno/NLT1: ntr
- Tiju/RAL: ntr
- Dimitri/KIT: ntr
- Zeeshan/NDGF: nta
- David/Dashboard:
  - ◆ comment on CMS problem, seems due to a VM crash, will investigate with SAM experts
  - ◆ can KIT get in touch with FTS developer to resolve the issue. [Dimitri/KIT: is there a ticket about this? David: will find out, otherwise will open one.]
- Massimo/Storage: nta

AOB: none

## Tuesday

Attendance: local (Andrea, Eric, David, Massimo, Eva, Alessandro, MariaD, Manuel); remote (Xavier/KIT, Kyle/OSG, Rolf/IN2P3, Burt/FNAL, Tiju/RAL, Jhen-Wei/ASGC, Ronald/NLT1, Giovanni/CNAF; Alessandro/ATLAS, Vladimir/LHCb).

Experiments round table:

- ATLAS reports -
  - ◆ SARA-MATRIX re-included in SantaClaus T0Export
  - ◆ ~09:00CET: from ATLAS Central Service team: "We are facing the problems with our IT virtual infrastructure. As far as I know some of IT hypervisors lost connection to SAN storage

Monday

devices. Some VM machines are blocked including three cassandra machines and one DDM machine. This was never happened in past. IT investigate it." There was a planned transparent intervention

<http://itsb.web.cern.ch/planned-intervention/service-consolidation-infrastructure-hardware-intervention>  
that was not transparent after all

<http://itsb.web.cern.ch/service-incident/cvicl04fc-cluster-failure/12-06-2012>

- ◆ CERN-PROD ALARM: ATLAS T0Merge files not accessible GGUS:83178 [↗](#) [Massimo: apologies for this issue, we need to change our procedures.]
  - ◆ FZK FTS Server channel stuck GGUS:83111 [↗](#) . Site Admin contacted FTS devs.
  - ◆ PIC FTS Server: proxy not valid? GGUS:83113 [↗](#) . Site Admin said that no action was taken, this could be related to one ongoing issue on FTS server with proxy renewal.
- CMS reports -
    - ◆ LHC machine / CMS detector
      - ◇ NTR
    - ◆ CERN / central services and T0
      - ◇ EOS space issues seem resolved and running smoothly
      - ◇ Numerous apparent false alarms this morning due to SLS issues. No issues since 10:00
      - ◇ [Massimo: noticed spikes of traffic in T1 transfers, is this normal? Eric: will investigate, thanks.]
    - ◆ Tier-1/2:
      - ◇ NTR
  - ALICE reports [↗](#) -
    - ◆ NTR
  - LHCb reports -
    - ◆ Users analysis and prompt reconstruction and stripping at T1s ongoing
    - ◆ MC production at Tiers2
    - ◆ T0:
      - ◇ CERN:
    - ◆ T1:
      - ◇ PIC : lhcbweb.pic.es does not respond (GGUS:83177 [↗](#)); Fixed

Sites / Services round table:

- Xavier/KIT: ntr
- Kyle/OSG: in maintenance period, some services may drop
- Rolf/IN2P3: ntr
- Burt/FNAL: minor degradation on tape services, has been resolved and should not have been visible
- Tiju/RAL: reminder, tomorrow will upgrade databases behind LFC/FTS and also will upgrade Castor the whole day
  - ◆ [Alessandro: how long for FTS? Tiju: scheduled 8am to 5pm but will be shorter than that, will take care of draining FTS channel one hour before. Alessandro: so whole UK cloud will be unavailable for transfers tomorrow; there is ICHEP conference and ATLAS is struggling to get MC production done. Tiju: this was discussed with ATLAS already.]
- Jhen-Wei/ASGC: ntr
- Ronald/NLT1: ntr
- Giovanni/CNAF: ntr
- David/Dashboard:
  - ◆ followup on FTS messaging from KIT, we are receiving messages, we just did not notice the endpoint had changed, sorry about that

- ◆ about CMS visualization there were two issues, one dashboard VM crash (will now get alarms to dashboard support list) and SAM ATP component issues (fixed in next release)
- Massimo/Storage: tomorrow network group 7-830 will intervene on switches affecting 10 our machines, conncted ATLAS T0 and they are ok
- Eva/Databases: ntr
- Manuel/Grid: ntr

AOB:

- MariaDZ: ATLAS ALARM ticket GGUS:82854 [↗](#) against T0 is still in status "Assigned" since 2012/06/04. [Massimo: the issue was fixed but we forgot to update the ticket, there was not a problem with the interface. MariaDZ: it is a pity because this incorrectly gives a bad statistics in the MB reports.]

## Wednesday

Attendance: local (Andrea, Luca, Eric, David, Manuel, MariaDZ); remote (Kyle/OSG, Jhen-Wei/ASGC, Ron/NLT1, Zeeshan/NDGF, Rolf/IN2P3, Gareth/RAL, Lorenzo/CNAF; Vladimir/LHCb, Alessandro/ATLAS).

Experiments round table:

- ATLAS reports -
  - ◆ RAL Oracle11 upgrade: whole UK cloud set to brokeroff
- CMS reports -
  - ◆ LHC machine / CMS detector
    - ◇ Access 1700-2200, physics ongoing
  - ◆ CERN / central services and T0
    - ◇ CASTOR disk server dropped out, restored, GGUS:83214 [↗](#).
  - ◆ Tier-1/2:
    - ◇ NTR
- ALICE reports [↗](#) -
  - ◆ NTR
- LHCb reports -
  - ◆ Users analysis and prompt reconstruction and stripping at T1s ongoing
  - ◆ MC production at Tiers2
  - ◆ T0:
    - ◇ CERN:
  - ◆ T1:
    - ◇ RAL : Scheduled Downtime
    - ◇ CNAF: Pilots failed; Fixed without GGUS ticket

Sites / Services round table:

- Kyle/OSG: ntr
- Jhen-Wei/ASGC: ntr
- Ron/ NLT1:
  - ◆ small network problem this morning for 1h, no GGUS ticket was sent so the impact on users cannot have been severe
  - ◆ dcache upgrade next Monday
- Zeeshan/NDGF: ntr

- Rolf/IN2P3: closed as unresolved the LHCb ticket about corrupted files, because data corruption does not seem specific to IN2P3 but rather seems a well known problem, for which the workaround is to use the checksum to check for data corruption and retry the transfer if necessary. Please look at GGUS:82247 [↗](#) for more details. [Vladimir: there are different reasons why there were corrupted files at IN2P3, the main one being site-specific problems in the transfer from WN to SE, but we agree that this ticket can be closed.]
- Gareth/RAL:
  - ◆ Castor upgrade to 2.1.11-9 was done and the services are ok now
  - ◆ problems were encountered during the 11g upgrade of the Oracle database behind FTS, a rollback had to be done, this is being followed up and ATLAS will be kept informed. [Alessandro: thanks for following up. Understood that the 11g upgrade at other sites was done using two instances, one on 10g and one on 11g, and only switching the alias at the end, was this done at RAL? Gareth: no we cannot do this as we only have one small RAC and this is the one being upgraded. We do have a small FTS instance that we could turn on in the meantime if the problem takes much longer, could this be useful? Alessandro: yes thanks, this could be useful. Will be followed up offline.]
- Lorenzo/CNAF: ntr
- Luca/Databases: ntr
- David/Dashboard: ntr
- Manuel/Grid: ntr

AOB:

- MariaDZ/GGUS:
  - ◆ two upcoming releases at unusual dates, Mon 25/6 with new reporting tools and Mon 9/7 with other development items in the pipeline
  - ◆ the BIOMED VO liked the team ticketing functionality and uses the TEAM ticket submit form for quite some time already. They wish to change the default ticket priority value from 'top priority' to a lesser value. If the experiments are happy with this request, please send feedback to Maria. Else the BIOMED request will be refused and we shall continue as now.

## Thursday

Attendance: local (Andrea, David, Maarten, Alessandro, Manuel, MariaDZ); remote (John/RAL, Zeeshan/NDGF, Rolf/IN2P3, Lisa/FNAL, Salvatore/CNAF, Jhen-Wei/ASGC, Ronald/NLT1, Gonzalo/PIC; Eric/CMS, Vladimir/LHCb).

Experiments round table:

- ATLAS reports -
  - ◆ CERN-PROD LSF bsub slow (response time always >1sec): ALARM GGUS:83252 [↗](#)  
[Manuel: this is a problem being experienced by all users, not just ATLAS. It is caused by high load, the traffic coming mainly from the CEs. Will change the master to a machine with a better network card tomorrow.]
  - ◆ RAL Oracle11 upgrade did not make it, site rolled back. To be re-scheduled. FTS server was returning for the jobs submitted before the start of the intervention exit code 256, which is the same of "no contact to FTS server", so ATLAS SS were polling all the already lost FTS jobs. Since this is definitely related to this aborted attempt of upgrade, we do not report further (i.e. no GGUS). [Maarten: status code 256 is not 8bits, probably this is status code 1 which is the default status code for failure, not a special status code for "no contact to FTS server". Alessandro: yes but the point is that the status code is the same and we use this value to decide about retrying, as we do see other status codes for other error situations. Maarten: open a GGUS ticket with FTS developers and ask to provide two different status codes for the two

cases. Massimo: may also be related to treatment of numbers by the new Oracle. Alessandro: thanks will open a GGUS ticket.]

- CMS reports -
  - ◆ LHC machine / CMS detector
    - ◇ NTR
  - ◆ CERN / central services and T0
    - ◇ LSF issues over the night. Failing and coming back up many times. Working OK as of this morning, but would like to understand if it's a load issue or related to recent upgrades.
    - ◇ VO virtual machine (vocms161) for CMSWEB lost a partition. Still not resolved. Suspect Hypervisor. Causing intermittent errors in crabserver dbs phedex reqmgr sitedb t0datasvc t0mon. Removed from our cluster. [Manuel: if you are seeing this problem since two days, it may be due to a hypervisor problem earlier this week, which can be fixed by a reboot. Maarten: please open a SNOW ticket and follow it up with the experts.]
  - ◆ Tier-1/2:
    - ◇ NTR
- ALICE reports [↗](#) -
  - ◆ NTR
- LHCb reports -
  - ◆ Users analysis and prompt reconstruction and stripping at T1s ongoing
  - ◆ MC production at Tiers2
  - ◆ T0:
    - ◇ CERN: RAW file not migrated (GGUS: 83239). [Vladimir: this is causing a backlog, what is the reason? Massimo: deep reasons not fully understood but there are several factors contributing to this: the age of the files is not being taken care of correctly; there is a high rate; the LHCb pool is too small. We are doing some configuration changes and we will also add new machines (also to replace the older machines that must be retired.)
  - ◆ T1:
    - ◇ ntr

Sites / Services round table:

- John/RAL: Oracle upgrade did not go very well, the DB is still down. Yesterday we decided to effectively prepare a new database for FTS. This is being used now but we will have to move back to the main database when the upgrade is completed.
- Zeeshan/NDGF: ntr
- Rolf/IN2P3: ntr
- Lisa/FNAL: ntr
- Salvatore/CNAF: ntr
- Jhen-Wei/ASGC: ntr
- Ronald/NLT1: ntr
- Gonzalo/PIC: ntr
- Kyle/OSG (after the meeting): ntr
  
- Manuel/Grid: nta
- David/Dashboard: ntr

AOB: (MariaDZ) Concerning the course in preparation at CERN on hadoop, announced last week, interested people from T1 centres can also attend but they will have to pay as any other course participant. Arranging

things via budget codes will be complex, so, we should know quite early how many people we are talking about. The course will take place mid-September at the earliest. Meanwhile, please send Maria candidate names.

## Friday

Attendance: local (Andrea, Maarten, Eric, Eva, Massimo, David); remote (Gonzalo/PIC, Ronald/NLT1, Xavier/KIT, Jhen-Wei/ASGC, John/RAL, Jeremy/GridPP, Rolf/IN2P3, Zeeshan/NDGF; Vladimir/LHCb, Doug/ATLAS).

Experiments round table:

- ATLAS reports -
  - ◆ CERN-PROD : overnight SRM transfer failure CERN-PROD\_TZERO, later in morning stages successful, ticket closed: GGUS:83292 [↗](#) [Massimo: not sure if ticket was a real issue or a glitch, anyway now looks ok. Doug: yes this seems to have been solved by itself.]
  - ◆ TRIUMF : FTS errors with proxy GGUS:83293 [↗](#)
- CMS reports -
  - ◆ LHC machine / CMS detector
    - ◇ Good data taking
  - ◆ CERN / central services and T0
    - ◇ NTR
  - ◆ Tier-1/2:
    - ◇ NTR
- ALICE reports [↗](#) -
  - ◆ NTR
- LHCb reports -
  - ◆ Users analysis and prompt reconstruction and stripping at T1s ongoing
  - ◆ MC production at Tiers2
  - ◆ T0: ntr
  - ◆ T1: ntr

Sites / Services round table:

- Gonzalo/PIC: ntr
- Ronald/NLT1: next Monday at SARA, tape lib maintenance and dcache upgrade
- Xavier/KIT: writing to tape was stalled today between 11am and 2pm
- Jhen-Wei/ASGC: ntr
- John/RAL: we were finally able to complete the DB upgrade. But the other database is still used for FTS and we still plan to switch back to the main FTS database.
- Jeremy/GridPP: ntr
- Rolf/IN2P3: ntr
- Zeeshan/NDGF: downtime next week on igf.si in Slovenia, some ATLAS files will be offline
- Eric/FNAL:
  - ◆ had a planned power outage earlier today; CMS was supposed to stay up but did not, we lost one computer room
  - ◆ planning a load shed to reduce risks with colling in the high Chicago temperatures, but we are still staying above the pledges
- David/Dashboard: ntr
- Massimo/Storage: nta



## WLCGDailyMeetingsWeek120611 < LCG < TWiki

- Manuel/Grid: batch situation at CERN is still being followed up. Interface CE/master was modified, resulting in reduced network traffic and improved response times, but there are still issues to be addressed, the situation is not yet as expected. The new machine with a better network card has also been identified but was not yet installed because it still needs to be cabled.
- Eva/Databases:
  - ◆ the security patches still need to be applied on the online DBs, this will be done during the technical stop in two weeks
  - ◆ during the technical stop we will also do some storage changes for all production databases, downtimes of 15 minutes are being planned with the experiments

AOB: none

-- JamieShiers - 22-May-2012

---

This topic: LCG > WLCGDailyMeetingsWeek120611

Topic revision: r18 - 2012-06-15 - 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