

Table of Contents

Week of 151019.....	1
WLCG Operations Call details.....	1
General Information.....	1
Monday.....	1
Thursday.....	3

Week of 151019

WLCG Operations Call details

- At CERN the meeting room is 513 [R-068](#).
- For remote participation we use the Vidyo system. Instructions can be found [here](#).

General Information

- The purpose of the meeting is:
 - ◆ to report significant operational issues (i.e. issues which can or did degrade experiment or site operations) which are ongoing or were resolved after the previous meeting;
 - ◆ to announce or schedule interventions at Tier-1 sites;
 - ◆ to inform about recent or upcoming changes in the experiment activities or systems having a visible impact on sites;
 - ◆ to provide important news about the middleware;
 - ◆ to communicate any other information considered interesting for WLCG operations.
- The meeting should run from 15:00 until 15:20, exceptionally to 15:30.
- The SCOD rota for the next few weeks is at [ScodRota](#)
- General information about the WLCG Service can be accessed from the [Operations Web](#)
- Whenever a particular topic needs to be discussed at the daily meeting requiring information from site or experiments, it is highly recommended to announce it by email to wlcg-operations@cernSPAMNOTNOSPAMPLEASE.ch to make sure that the relevant parties have the time to collect the required information or invite the right people at the meeting.

Monday

Attendance:

- local: Herve (storage), Maarten (SCOD + ALICE)
- remote: Alessandro (ATLAS), Asa (ASGC), Christian (NDGF), Di (TRIUMF), Francesco (CNAF), Gareth (RAL), Lisa (FNAL), Michael (BNL), Oliver (CMS), Onno (NLT1), Pepe (PIC), Rolf (IN2P3), Sang-Un (KISTI), Stefano (LHCb)

Experiments round table:

- ATLAS reports (raw view) -
 - ◆ FTS "stalled connection": now problem manifested itself also with RAL instance. FTS developers suggested a patch, which is now being applied.
 - ◆ Tier-0 TZDISK number of inode issue, ALARM ticket created and immediately this number was increased.
- CMS reports (raw view) -
 - ◆ 25ns Data taking, normal operation
 - ◆ Production activity slowly tailing off
 - ◆ CERN:
 - ◇ GGUS:117001 [?](#) : All SAM pilot tests for all CEs for CERN only in (same) error. Maybe the same as in GGUS:116468 [?](#)? Argus?
 - Maarten: *being looked into*
- ALICE -

- ◆ high activity
- LHCb reports (raw view) -
 - ◆ Data Processing:
 - ◇ Data processing of pp data at T0/1/ sites, Monte Carlo mostly at T2, user analysis at T0/1/2D sites
 - ◇ Stable number of running jobs for processing of data at T0/1
 - ◇ Data from pHe fully processed.
 - ◆ T0
 - ◇ Ticket about FTS409 stalled: solved reporting some log (GGUS:116965 [↗](#))
 - ◇ Ticket about aborted SRM at CERN (GGUS:116897 [↗](#)) and (GGUS:116877 [↗](#)) should be closed. Wait for some more data transfer from the pit before confirming the problem is solved.
 - ◇ *Maarten: there also was the alarm ticket GGUS:116973 [↗](#) that was opened on Sun because of a problem with EOS*
 - ◆ T1
 - ◇ SARA storage is in scheduled downtime
 - ◆ T2
 - ◇ FTS transfer failing at UKI-SOUTHGRID-RALPP (GGUS:116995 [↗](#)). Investigation ongoing.

Sites / Services round table:

- ASGC: ntr
- BNL: ntr
- CNAF:
 - ◆ Some problems on ATLAS disk (one ddl currently offline), 350 TB data are temporary not available. More news as soon as possible (we are still performing some checks)
- FNAL: ntr
- GridPP:
- IN2P3: ntr
- JINR:
- KISTI: ntr
- KIT:
- NDGF:
 - ◆ A tape robot lost its arm during the night. An IBM technician arrived on site recently and the spare parts will arrive at 17:30 CEST. For now the disks are keeping up with the writes, and the whole thing should be transparent if we get it fixed tonight. There might be some long load times for Atlas and Alice data.
 - ◆ Tomorrow there will be a short downtime as dCache pools are updated. Only short interruptions during reboots. Some Atlas and Alice data unavailable.
 - ◆ Both are registered in gocDB.
- NL-T1:
 - ◆ SARA downtime for network maintenance being used also for moving some dCache components to cure space manager timeout errors
 - ◆ Vidyo first broke our connection to this meeting, then connected us to another room; on the 3rd attempt we were back in the right room
 - ◇ *Maarten: we will report this to the Vidyo team*
- NRC-KI:
- OSG:
 - ◆ the bad version of JGlobus that broke BeStMan and EOS installations last week has been replaced with the correct rpm, so it should again be safe to upgrade
- PIC:
 - ◆ we uploaded a revised version of the Nagios Plugin for SAM to github. Some bugs were

fixed, the rpm can be found here:

◇ <https://github.com/jcasals/nagios-plugins-lcgsam>

- ◆ we are planning a downtime for the 24th of November to upgrade dCache to the latest golden release. Is this ok for the experiments?
- RAL:
 - ◆ during the weekend there were some problems with disk servers used by ATLAS and LHCb, who were subsequently notified
- TRIUMF: ntr

- CERN batch and grid services:
- CERN storage services:
 - ◆ Nothing to report
- Databases:
- GGUS:
 - ◆ GGUS was temporarily unavailable Mon afternoon due to a power cut at KIT
- Grid Monitoring:
- MW Officer:

AOB:

Thursday

Attendance:

- local: Asa (ASGC), Eric (CMS), Herve (storage), Iain (batch & grid), Maarten (SCOD + ALICE), Stefano (LHCb), Zbyszek (databases)
- remote: Christian (NDGF), Christoph (CMS), Dario (ATLAS), Dennis (NLT1), Di (TRIUMF), Gareth (RAL), Kyle (OSG), Matteo (CNAF), Rolf (IN2P3), Thomas (KIT)

Experiments round table:

- ATLAS reports (raw view) -
 - ◆ It is not clear which version of FTS is installed at each site. Several patches came and went and we don't know which site has which patch to correlate to the problems we see.
 - ◇ *Maarten: we will discuss that in the Ops Coordination meeting later today*
- CMS reports (raw view) -
 - ◆ Production system is full since Tuesday. 100K running production jobs. 20K running analysis jobs
 - ◆ Taking data as available
- ALICE -
 - ◆ NTR
- LHCb reports (raw view) -
 - ◆ Data Processing:
 - ◇ Data processing of pp data at T0/1/2 sites. Some T2 attached to T1 in order to speed up the processing.
 - ◇ Monte Carlo mostly at T2, user analysis at T0/1/2D sites
 - ◇ Stable number of running jobs for processing of data at T0/1.
 - ◇ Data from pHe fully processed.
 - ◆ T0
 - ◇ Transfer to EOS is now OK. No more failures observed since Monday.

- ◇ Problems related to the transfer of large number of small files from pit has been solved putting in place a merging procedure before transfer. Working since yesterday and no obvious problem observed so far.
- ◇ A ticket regarding connections to fts407 hanging has been submitted just after lunch (GGUS:117128 [↗](#)).

◆ T1

- ◇ Few raw files lost at RAL due to a diskserver went down during the weekend. Files are being re-replicated.

Sites / Services round table:

- ASGC: ntr
- BNL:
- CNAF:
 - ◆ still investigating the issue affecting some of the ATLAS disks
- FNAL:
- GridPP:
- IN2P3: ntr
- JINR:
- KISTI:
- KIT:
 - ◆ Mon afternoon a campus router went down when a power cable was hit by a digger
 - ◇ GGUS was unreachable for ~2h
 - ◇ LHCOPN and LHCONE were not affected
- NDGF:
 - ◆ the issues mentioned on Mon have been fixed
 - ◆ the reverse DNS for the IPv6 address of srm.ndgf.org has also been fixed
 - ◇ that service should be fully IPv6-compliant now
- NL-T1: ntr
- NRC-KI:
- OSG: ntr
- PIC:
- RAL: ntr
- TRIUMF: ntr
- CERN batch and grid services:
 - ◆ HTCondor Grid pool becomes production on the 2nd of November.
 - ◆ Beginning 02/11/2015 we will begin the official decommissioning of the ARC CEs leading to our HTCondor Pool.
 - ◆ Leaving just one CE type, the HTCondorCE. The decommission phase will be lengthened by request of a VO, to give them extra time.
 - ◆ *Maarten: progress was made on resolving the recurrent failures of the CMS pilot role SAM tests*
 - ◇ *the number of accounts in the CMS pilot pool is rather low*
 - ◇ *this has led to frequent recycling of such accounts*
 - ◇ *that in turn implied an increased exposure to a race condition in Argus*
 - ◇ *ticket GGUS:117125 [↗](#) has been opened for the Argus devs to look into that*
 - ◇ *meanwhile the number of such accounts will simply be increased*
 - *that is desirable anyway for improved traceability*
- CERN storage services:
 - ◆ EOSALICE Update finished at ~14:30, took longer than expected, the maintenance window has been extended to fix a locking bug
- Databases:
 - ◆ yesterday the read-only replica of the ATLAS ADCR DB was down from ~12:30 to ~13:00

due to a crash

◇ this affected Rucio and PanDA monitoring

◇ the cause was an Oracle bug for which a fix will be tested in the next days

- GGUS:
- Grid Monitoring:
- MW Officer:

AOB:

This topic: LCG > WLCGDailyMeetingsWeek151019

Topic revision: r13 - 2015-10-22 - MaartenLitmaath



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](#)