

# Table of Contents

<b>WLCG Operations Coordination Minutes - June 19th, 2014.....</b>	<b>1</b>
Agenda.....	1
Attendance.....	1
General News.....	1
Middleware News.....	1
Oracle Deployment.....	2
Tier 0 News.....	2
Tier 1 Feedback.....	2
Tier 2 Feedback.....	2
Experiments Reports.....	2
ALICE.....	3
ATLAS.....	3
CMS.....	3
LHCb.....	4
Ongoing Task Forces and Working Groups.....	4
Tracking Tools Evolution TF.....	4
FTS3 Deployment TF.....	4
gLExec Deployment TF.....	4
Machine/Job Features.....	5
Middleware Readiness WG.....	5
Multicore Deployment.....	5
SHA-2 Migration TF.....	6
WMS Decommissioning TF.....	6
IPv6 Validation and Deployment TF.....	6
HTTP Proxy Discovery TF.....	6
Network and Transfer Metrics WG.....	6
Action list.....	7
AOB.....	7

# WLCG Operations Coordination Minutes - June 19th, 2014

## Agenda

- <https://indico.cern.ch/event/323689/>

## Attendance

- Local: Maria Alandes Pradillo (minutes), Marian Babik, Simone Campana, Andrej Filipcic, Alessandro di Girolamo (ATLAS), Maarten Litmaath (ALICE), Oliver Keeble, Felix Lee (ASGC), Andrew McNab (LHCb), Hassen Riahi, Stefan Roiser.
- Remote: Maite Barroso, Frederique Chollet (IN2P3), Michael Ernst, Alessandra Forti (chair), Isidoro Gonzalez Caballero, Burt Holzman (FERMILAB), Michel Jouvin, John Kelly, Peter Love, Antonio Perez-Calero Yzquierdo, Di Qing (TRIUMF), Rob Quick (OSG), Gareth Smith, Chris Walker.

## General News

- Meeting in 2 weeks canceled because there is the WLCG Workshop in Barcelona.
- Write reports on a TF page and use INCLUDE variable for the WLCG Ops Coordination report. You can look at the multicore report as an example. The report is included from Multicore TF Report now.

## Middleware News

- Useful Links:
  - ◆ Baseline Versions
  - ◆ MW Issues
  - ◆ Storage Deployment
  - ◆ Data Management news
- Baselines:
  - ◆ Storm 1.11.4 released in EMI containing several bug fixes. Still 1.11.3 as baseline, waiting for the UMD release.
- MW Issues:
  - ◆ Issue "Job submission fails for VOs supported by VOMS server with SHA-512 host certificate": the fix for bouncycastle-mail was released both by EMI and UMD. The SHA2 TF should coordinate the upgrade in prod of the affected services ( CREAM, UI, WN and Argus)
  - ◆ 3 issues affected some sites after the latest EMI update of Cream and LB. The problems are under investigations by the PTs, for the moment the affected components have been removed from the EMI repo and the sites that have already upgraded had to downgrade to the previous version of the software and they are working fine now.
- T0 and T1 services
  - ◆ No short term planned intervention
  - ◆ Recent changes:
    - ◆ BNL :
      - ◇ ATLAS LFC @ decommissioned following the migration to Rucio
      - ◇ FTS3 production instance deployed and tested
    - ◆ CERN
      - ◇ Castor upgrade to 2.1.14-14 ( latest)
      - ◇ EOS updates to 0.3.35/xrootd-3.3.6 (plus patches) ( latest)
    - ◆ KIT

- ◊ dCache upgrade to 2.6.29 (latest) for CMS
- ◆ PIC
  - ◊ dCache upgrade to 2.6.29 (latest)
- ◆ NL-T1 ( SURFSARA)
  - ◊ dCache upgrade to 2.6.29 (latest)
- CVMFS upgrade to 2.1.19
  - ◆ Broadcast sent last week to all WLCG sites ( CVMFS\_broadcast\_mail.txt ) ( please ask to the site managers connected if they all received it)
  - ◆ Sites running CVMFS 2.1.19 ( total includes also sites where CVMFS is not installed or is not monitored)
    - ◊ ALICE: 27/100
    - ◊ LHCb: 32/87
    - ◊ CMS : 41/131
    - ◊ ATLAS: Still SAM probe not configured, A. di Girolamo is working on this.
  - ◆ Starting from July, the sites not compliant with the 2.1.19 version will be notified with a GGUS ticket.

Alessandra and Maarten mention that the CVMFS upgrade just requires an update of the RPM, so it should be very easy for sites to apply this upgrade.

## Oracle Deployment

- Upgrade Plans Table

## Tier 0 News

- lxbatch and lxbplus upgraded with CVMFS client  $\geq 2.1.19$ , as requested at last meeting
- Migration to VOMS-admin: new release fixing GGUS:102984 [↗](#) was expected 1-2 weeks ago, no news yet from VOMS-admin team
  - ◆ The OPS VO now runs in voms-admin instead of VOMRS, after the migration done on June 17th
- The SLC5 submitting CEs (ce201-ce207) were put into draining mode today; no new submissions are allowed. With this we close the SLC5 to SLC6 migration reporting in WLCG, only local submissions to SLC5 are allowed, and we are gradually migrating the remaining capacity.
- Job efficiency, recent meeting on June 13th:
  - ◆ live efficiency dashboard [↗](#) available
    - ◊ Alessandro will run sample Hammercloud jobs on regular basis and provide jobids to plot in the dashboard (otherwise analyzing the dashboard plots is very difficult, as all jobs are mixed, and some of them are not efficient by nature)
  - ◆ Next meeting on 1 month

Alessandra mentions that Maite could then close the GGUS ticket tracking the migration to SLC6.

## Tier 1 Feedback

## Tier 2 Feedback

## Experiments Reports

## ALICE

- successful campaign for users to move away from old ROOT versions
  - ◆ user analysis jobs now read data locally when the local SE is in good shape
- CERN
  - ◆ SLC6 job efficiencies:
    - ◇ various data analytics and comparison efforts ongoing
    - ◇ Meyrin job efficiencies look OK
    - ◇ Wigner jobs (both on VM and HW) appear to have significantly lower efficiencies
    - ◇ live plot: <http://alimonitor.cern.ch/?1201>

Maite mentions that ALICE conclusion on SLC6 job inefficiencies is not aligned with the work that is being carried out by the IT group of experts following up this issue. Maite adds that this problem needs to be understood better before reaching any conclusion. Alessandro also adds that ATLAS is not reporting anything in particular for this issue because ATLAS agrees that this is centralised by Maite's report which they fully agree with.

## ATLAS

- Rucio full chain test progressing. Discovered, "thanks" to the CERN FTS3 issues of the week of the 10th of June which created a huge backlog in data transfers, a race conditions between conveyor (the agent which submits and poll the data transfers via FTS3 jobs) and the reaper (the agent which take care about deleting the data once the lifetime has expired). Test will keep on for at least other 3 weeks
- DC14 expected to start in approximately 2 weeks from now. The sites should not actually see any difference respect to the present activity.
- Reprocessing of Heavy Ion data10 has started this week. We remind Tier1s that the recall from tape is automatically managed by Panda so in theory there should be no manual operations to be done by the Tier1s .
- Panda/Jedi is now fully ready for user analysis: the missing bit was the monitoring which is now available.
- the CERN FTS3 issue of the past week created a backlog in the production data transfers which was reflected in approx two day of part of the Grid draining: while recovering the backlog, we observed that the whole DDM system wasn't able to transfer more than approx 2Mfile a day: we are now investigating the limiting factors.

Alessandro clarifies that the problems due to the FTS3 issue make things to be very unstable and trigger other problems as the backlog is huge and is not digested on time, creating an avalanche effect due to the high number of active jobs. They expect normal data flow again after the workaround applied to Rucio/DQ2.

## CMS

- updated site readiness
  - ◆ after introduction of disk/tape separation at the T1 sites, we needed to adapt site readiness for T1 sites
    - ◇ previously, we checked that a T1 site had enough active PhEDEx links (debug traffic runs without problem) and good PhEDEx links (quality of production transfers is good) only for the tape PhEDEx endpoint
    - ◇ now we have the same expectations for the disk and the tape PhEDEx endpoint, meaning that if disk or tape endpoint are failing the criterias, the site is not ready
  - ◆ We also reactivated the requirements on the PhEDEx links from the T0, data taking is coming
- Russian Tier-1 T1\_RU\_JINR
  - ◆ Used -successfully- for MC so far
  - ◆ Extending to re-processing workflows
- Downtime of GGUS last weekend

- ◆ Noticed by CMS, but no big impact
- ◆ Fixed by experts even during the weekend
- CMS SAM tests
  - ◆ All major sites are passing
  - ◆ Plan to make xrootd-fallback test critical: June 30th, 2014
- Started to remove individual release tags from CEs
  - ◆ Not used by recent CMS submission tools
  - ◆ Will unload the BDII significantly
  - ◆ Good fraction already removed

## LHCb

- Operations
  - ◆ Mainly Monte Carlo and user jobs executed since the last meeting
  - ◆ Test reprocessing of 2010 data carried out and being validated before full campaign
- CVMFS recommended (== new baseline) client version 2.1.19 installed by 33 sites so far (Up from ~20 last time)
  - ◆ Remaining sites are mostly on 2.1.17 (15) or 2.1.15 (11), 2.1.14 (4) -> this may be risky
  - ◆ See dashboard view [↗](#) for site-by-site versions
- We ask sites to ensure that downtimes, including unscheduled outages, accurately reflect the specific services which are unavailable. We use this information in the RSS to automatically stop sending sending jobs/transfers to sites and it requires manual intervention if the downtimes aren't accurate.

Alessandra mentions that most of the sites should be declaring GOCDB downtimes properly as there are other VOs, like ATLAS, also relying on this. Maarten asks whether they have followed with up with the concerned sites using GGUS. Andrew explains that this has indeed been followed up with GGUS, but that LHCb considered useful to make a general reminder in the meeting.

## Ongoing Task Forces and Working Groups

### Tracking Tools Evolution TF

#### FTS3 Deployment TF

- Monitoring the auto-tuning algo. closely and adjusting various monitoring tools of FTS3
- CMS
  - ◆ ASO will start to use this week FTS3@RAL instead of FTS3@pilot for CSA14
  - ◆ opening tickets to complete also migration of PhEDEx Prod transfers to FTS3
- Last Monday (9th of June) all FTS3 servers at CERN encountered random disconnections from MySQL database. Is it still not understood whether it was networking or database related issue and, as a result, the service was unstable for 24h

#### gLExec Deployment TF

- 85 tickets closed and verified, 10 still open (no change)
- Deployment tracking page
- other activities on hold until resolution of Argus stability issues
  - ◆ e.g. GGUS:105666 [↗](#)
  - ◆ the Argus support situation is still unclear

Maite clarifies that CERN incident mentioned by Maarten in the minutes was not due to ARGUS. Maarten changes the minutes accordingly. Alessandra asks what the current situation with ARGUS support is. Maria explains that Andrea Sciaba is in contact with INFN waiting for more details. There are no news yet. Michel

adds that WLCG management is aware of this.

Alessandra asks whether there is something to be done for the missing sites not deploying gLExec. Maarten answers that nothing will be done for the time being.

## Machine/Job Features

- Batch system status
  - ◆ PBS/TORQUE implementation: Deployment done at NIKHEF, UB-LCG2 will try to deploy as second test site
  - ◆ LSF: CNAF implementation done for machinefeatures, jobfeatures need a bit more tweaking
  - ◆ SLURM implementation pending
- Project Status

After request from Alessandra, Stefan gives details on SGE and HTCondor implementations, where there are also a couple of sites deploying them.

## Middleware Readiness WG

- DPM: progress on setups for ATLAS and CMS, and on preparation of 1.8.9 release
  - ◆ see tasks overview
- HTCondor: contact and testing details received from OSG
  - ◆ added to product teams table
- Tracking of MW versions deployed at sites:
  - ◆ proposal discussed in June GDB and afterward, not yet converged
  - ◆ prototype being implemented, to be presented at the July 2 meeting of the WG

Alessandra asks when the first version of the prototype will be ready. Maarten explains that this will be presented at the next WG meeting. Michel asks whether there has been any conclusion for deciding not to use Pakiti, Maarten explains Michel that Lionel has all the technical details and that it is better to follow up with him. Again all the details will be presented and discussed at the next WG meeting.

## Multicore Deployment

- CMS has been running multicore pilots for single core production jobs at most CMS T1s (PIC, KIT, RAL, JINR and CCIN2P3) for over a month now. Since we are sending a mix of single core and multicore pilots to pull jobs from the same pool, whenever there is workload for the site (which is most of the time recently), multicore pilots are sent. So this means actually that the sites are constantly receiving a relatively stable amount of multicore pilots from CMS.
- With respect to feedback from the sites regarding this activity, we don't have a detailed report from any of them yet, but no complaints either so far. We are expecting to collect reports from a few sites (for example KIT, RAL and PIC) in a coming Multicore TF meeting in order to be ready for the coming WLCG workshop.
- ATLAS hasn't restarted multicore submission yet due to problematic software validation.

ATLAS clarifies that they are waiting for a new version of the software, instead of software validation as mentioned in the TF report.

*Alessandra post meeting clarification: The new releases are currently going through validation.*

Antonio adds that the TF is now waiting for the sites reports and at the same time is also working on monitoring tools to be able to understand better job efficiency, run time, etc metrics. Antonio also clarifies that CMS schedules a mix of both single and multicore pilots.

## SHA-2 Migration TF

- introduction of the new VOMS servers
  - ◆ the fix for blocking issue [GGUS:104768](#) is part of [UMD 3.7.0](#) released Thu Jun 12
    - ◇ see list of known issues
    - ◇ [bouncycastle-mail-1.46-2](#)
    - ◇ affected node types: ARGUS, CREAM, UI, WN
  - ◆ all sites need to update their affected hosts
  - ◆ experiments have to update their affected *UI* instances, if any
  - ◆ a broadcast will be done when a new timeline has been agreed
- RFC proxies
  - ◆ no progress on the open issues due to other priorities

Maarten clarifies that a new deadline is being discussed with VOMS managers after the fix has been released in UMD. Broadcasts have been sent to the sites to ask them to upgrade their nodes.

## WMS Decommissioning TF

- Very good progress in SAM Condor validation since the last meeting. Overall, there were 12 issues identified so far, 6 were resolved. Among those resolved, the most critical one was Condor HELD status for many US sites with reason:
  - ◆ "Globus error 155: the job manager could not stage out a file" - fixed in Condor job submission probe by adding outputs as inputs in JDL
- Issues we're currently working on:
  - ◆ Condor HELD status due to: CREAM\_Delegate Error: Connection to service ....  
FaultDetail=[SSL authentication failed in tcp\_connect(): check password, key file, and ca file.  
◇ Investigation ongoing, impacted testing of many ATLAS sites on Tuesday
  - ◆ Condor GAHP core dumps - we have ~ 1000 seg. faults of cream\_gahp/day in CMS, ~ 600/day in ATLAS - there seems to be no major impact on measurements, meaning SAM Condor probe can recover from the errors
  - ◆ Sporadic errors for some sites, e.g. Globus error 121: the job state file doesn't exist, Unspecified grid manager error
  - ◆ ARC-CE WN tests failing for CMS - impacts just 3 sites, but all their CEs (T2\_FI\_HIP, T2\_EE\_Estonia, T2\_UK\_London\_IC)
    - ◇ Proxy doesn't exist on the worker node (Error: X509\_USER\_PROXY points to a non existing location)

Andrej mentions that ATLAS can have a look at the ARC CE issues. Maarten mentions that the Condor team is aware of the seg fault issues and that this seems to be a real bug.

## IPv6 Validation and Deployment TF

### HTTP Proxy Discovery TF

- Task table was created

## Network and Transfer Metrics WG

- Invitations were sent to all cloud contacts, very good feedback received. Roles and responsibilities documented at <https://twiki.cern.ch/twiki/bin/view/LCG/MeshLeaders>
- Planning to organize a kick-off meeting in the second or third week in July
- Scheduled meeting on OSG Network Datastore to discuss current status and plans (<https://twiki.grid.iu.edu/bin/view/Production/OSGNetworkDatastorePlan>)

## Action list

1. ONGOING on the WLCG monitoring team: status of the CondorG probes for SAM to be able to decommission SAM WMS
2. ONGOING on the middleware officer: report about progress in CVMFS 2.1.19 client deployment
3. NEW on the WLCG monitoring team: evaluate whether SAM works fine with HTCondor CE. Report about showstoppers.

## AOB

Maarten explains that OSG has reported at the 3 o'clock meeting that they are planning to migrate to HTCondor CEs in the autumn and that this would require the SAM tests to work properly with this new CE type. Rob confirms that the plans are to recommend the default CE installation to be HTCondor by November 1st 2014. GRAM 5 CEs will be supported for at least mid to late 2015. Maarten explains that maybe this deadline will not be possible if there are problems with SAM. Rob says that OSG will adjust it as necessary. Marian mentions that in principle he doesn't see any showstoppers but that this will have to be tested. Maarten suggests that the SAM team tries the HTCondor CE and reports back any possible issues. Maria proposes to have an action item on this so that WLCG monitoring can report in the upcoming weeks.

-- MariaALANDESPRADILLO - 16 Jun 2014

---

This topic: LCG > WLCGOpsMinutes140619

Topic revision: r36 - 2018-02-28 - 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