

# Table of Contents

<b>WLCG Operations Coordination Minutes, November 5th 2015.....</b>	<b>1</b>
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Middleware News.....	1
Tier 0 News.....	2
DB News.....	2
Tier 1 Feedback.....	2
Tier 2 Feedback.....	2
Experiments Reports.....	2
ALICE.....	2
ATLAS.....	2
CMS.....	3
LHCb.....	4
Ongoing Task Forces and Working Groups.....	4
gLExec Deployment TF.....	4
HTTP Deployment TF.....	4
Information System Evolution.....	4
IPv6 Validation and Deployment TF.....	5
Middleware Readiness WG.....	5
Multicore Deployment.....	5
Network and Transfer Metrics WG.....	5
RFC proxies.....	6
Squid Monitoring and HTTP Proxy Discovery TFs.....	6
Action list.....	6
Specific actions for experiments.....	7
Specific actions for sites.....	7
AOB.....	7

# WLCG Operations Coordination Minutes, November 5th 2015

## Highlights

- CERN and OSG developed a BDII info provider for HTCondor CE, already running at CERN. The package will be distributed via the WLCG repository so other sites running HTCondor CE can benefit from it.
- perfSONAR collector, datastore, publisher and dashboard now in production.

## Agenda

- <https://indico.cern.ch/event/393619/>

## Attendance

- local: Maria Alandes (Minutes, Infosys TF), Maarten Litmaath (ALICE), Andrea Sciaba (IPv6 TF), Andrea Manzi (MW Officer), Jerome Belleman (T0), Marian Babik (Network WG), Sabine Crepe-Renaudin, Katarzyna Maria Dzedziniewicz-Wojcik (IT-DB), Alessandro di Girolamo (ATLAS)
- remote: Alessandra Forti (Chair), Michael Ernst, Christoph Wissing (CMS), David Cameron, David Mason, Gareth Smith, Felix Lee, Antonio Maria Perez Yzquierdo.
- apologies: Catherine Biscarat

## Operations News

- Stefan has stepped down as Machine/Job Features task force leader. Until a new TF coordinator is found, you can keep on contacting the TF through their mailing list ([wlcg-ops-coord-tf-machinejobfeatures@cernNOSPAMPLEASE.ch](mailto:wlcg-ops-coord-tf-machinejobfeatures@cernNOSPAMPLEASE.ch)) for any matter concerning the TF activities. Volunteers for taking on the role are very welcome.

## Middleware News

- Useful Links:
  - ◆ Baseline Versions
  - ◆ MW Issues
  - ◆ Storage Deployment
- Baselines:
  - ◆ CERN and OSG developed a BDII info provider for HTCondor CE, already running at CERN. The package will be distributed via the WLCG repository so other sites running HTCondor CE can benefit.
- Issues:
  - ◆ Issue affecting TopBDII and Resource BDII in ARC-CE with the latest version of openldap (2.4.40-5/6) on SL6/CentOS6. We have still not updates from RedHat regarding a fix.
  - ◆ FTS access via REST API was affected by a problem with stalled connections . FTS site managers have been asked to configure Apache to use mpm\_worker mode and everything looks ok now.
  - ◆ Broadcast by EGI SVG to sites:

◇ Critical java vulnerabilities affecting java 6,7,8

<https://wiki.ege.eu/wiki/SVG:Advisory-SVG-2015-9707> : sites are asked to updated to the latest openjdk for 1.7 and 1.8, but no fix is going to be provided for openjdk 1.6. Still some grid components depend on openjdk 1.6 ( like WN) so we need to understand if it would be possible to move to openjdk 1.7 without problems.

Discussions ongoing

- T0 and T1 services
  - ◆ CERN
    - ◇ EOS Alice updated to v 0.3.134-aquamarine
  - ◆ JINR
    - ◇ dCache updated to v 2.10.42
  - ◆ NDGF
    - ◇ dCache updated to v 2.13.12

## Tier 0 News

- HTCondor grid service in production since 2 Nov.
- LSF 9 master upgrade on 4 Nov successful.

## DB News

IT-DB is preparing the installation of October Critical Patch and OS updates in the test and integration databases. This will happen in the coming weeks. All interventions are planned to be rolling. The upgrade of production databases will take place in January.

## Tier 1 Feedback

## Tier 2 Feedback

## Experiments Reports

### ALICE

- generally high activity
  - ◆ 90k+ for a few hours on Oct 30
  - ◆ Reco, analysis, MC
- CERN and Tier-1 sites that will receive heavy ion data have been asked for high memory queues:
  - ◆ to be used only for raw data reconstruction
  - ◆ until early next year (~Feb)

### ATLAS

- Activity as usual:
  - ◆ running around 200k jobs. No special issues.
  - ◆ spillover of Tier-0 data now successfully working. Now running "in parallel mode", i.e. on data the Tier-0 already processed, experiencing massively (400M events). Order of 25k slots running. We will work in this mode for the rest of p-p data taking to perform extensive checks and optimizations.
  - ◆ mc15b, 0.5B digi+reco is starting now. It can be heavy on sites in terms of I/O.
  - ◆ reprocessing of Run2 data: foreseen to start mid-December
- FTS bringonline : BNL experts noticed issues on FTS server related to bringonline requests. These

issues were slowing down considerably the whole FTS server (at BNL). Several interaction between FTS BNL service manager (Hiro) and FTS devels experts. It seems that now the problem is solved. It also seem that now FTS in BNL is going faster (still to be confirmed): maybe the changes applied can be beneficial for all the FTS servers. FTS devels need to comment.

- TAPE: we are having reports from few sites (e.g. RAL, IN2P3 -CC, FZK) that ATLAS is using the tapes "very heavily". Sometimes the problem is related to the increase usage of tape wrt the present buffer disk size, but also we noticed that over the same month same file was recalled few times (i.e. there is room for optimization on workflow) . Actions taken this week: we reduced the weight used by Panda to broker to sites with tape replica if any other disk replica available, and we are now observing decrease of usage of tapes from this point of view (which is small in terms of load, but as we said it was some workflow optimization). Other point TAPE related: During the TS+MD we will consolidate 1PB of "old not heavily accessed data" on tapes (for info we write usually 2PB/months, but without RAW data the throughput should be similar to the present one).
- TAIWAN-LCG2 Tape fully decommissioned.
- Consistency checks. Storage dumps: we are asking the site to provide the storage dumps. please check <http://go.web.cern.ch/go/C9xr> (for ATLAS people [https://indico.cern.ch/event/445782/contribution/13/attachments/1180967/1709690/proposal\\_to\\_sites.pdf](https://indico.cern.ch/event/445782/contribution/13/attachments/1180967/1709690/proposal_to_sites.pdf) )
- Monitoring: we need a reliable and meaningful monitoring for central service => CVMFS stay in orange in Kibana ([https://meter.cern.ch/public/\\_plugin/kibana/#/dashboard/temp/ATLAS::ADC\\_CS](https://meter.cern.ch/public/_plugin/kibana/#/dashboard/temp/ATLAS::ADC_CS)), what is monitored there ? how to improve ?

Andrea Manzi explains that there will be a new release of FTS where bringonline requests will be limited to 1K. This limit will imply less stress in the DB.

Following a question by Maarten Litmaath, Alessandro di Girolamo confirms that the monitoring issue is being followed up with the right people.

## CMS

- Good resource utilization in October
  - ◆ Frequently 150k cores filled
  - ◆ Reached a peak of 200k+ parallel jobs on Oct 9th
  - ◆ Erratum - last report
    - ◇ Quoted 120k parallel jobs were counting cores by central production only
- Change in DDM policy
  - ◆ Datasets that are on tape and are not accessed over 18 months will be removed form disk
  - ◆ Getting data back to disk requires operator intervention for now
- Pilot flavours sent to Tier-1 sites
  - ◆ Present approach
    - ◇ Pilots with VOMS role 'production' (~90%) of CMS share
    - ◇ Pilots with VOMS role 'pilot' (~10) of CMS share
    - ◇ Sending two types limits scaling
  - ◆ Two scenarios with one flavour
    - ◇ Only 'production' pilots
      - Requires no change in fairshare settings at sites
      - Pilot might not be allowed to use glexec, needed for analysis jobs
    - ◇ Only 'pilot' pilots
      - Requires likely changes in fairshare settings
      - No changes in glexec configuration
    - ◇ Should be decided soon
- Campaign to re-new Phedex credentials went fine
  - ◆ A few Tier-2 sites did not manage to succeed in 48 hours

- ◆ Still unconnected sites are receiving GGUS tickets
- T2\_RU\_RRC\_KI will drop Tier-2 support for CMS: GGUS:117327 [↗](#)
  - ◆ Will be removed from CMS systems soon

Antonio explains that CMS will switch temporarily to running all jobs, production and analysis, at T1s via a single type of pilot jobs. CMS expects to have more flexibility in job scheduling prioritisation and improve performance in terms of CPU utilization. At the moment, CMS submits two types of pilot jobs to T1s: analysis and production. Relative workload prioritization at T1s is achieved indirectly by local prioritization of the pilots. Production pilots have more priority at T1s, but they can't run analysis jobs, hence, If the production load decreases, resources allocated to production pilots may be underutilised. With one type of pilot jobs, the same resources will be more easily re-allocated from one type of workload to the other. If the testing phase works as expected, CMS will propose to make the change permanent.

## LHCb

- Operations
  - ◆ Very high activities on distributed computing resources with data processing and simulation workflows, including pit export and T0 replication of RAW data
  - ◆ Some T2 sites have been included into the data processing workflows
  - ◆ LHCb will participate and take data in led-ion runs until mid December
- Issues
  - ◆ Several CERN/EOS failures with data access (GGUS:117368 [↗](#), GGUS:117302 [↗](#), GGUS:117157 [↗](#)). Many thanks to EOS team to always quickly react and get the service running (mostly restart of bestman)
  - ◆ Two problems where CERN/FTS became unresponsive (GGUS:117206 [↗](#), GGUS:117128 [↗](#)), also quickly fixed, many thanks
- Development / Outlook
  - ◆ Working on interface to HTCondor-CE
  - ◆ Workflows with DBC clouds have been tested successfully, ready for further submissions with monte carlo simulation

## Ongoing Task Forces and Working Groups

### gLExec Deployment TF

- NTR

### HTTP Deployment TF

### Information System Evolution

- Input for Future Use Case document is being finished within the experiments, some drafts already available waiting for the final green light. A complete first draft will be distributed within the TF in the next days.
- Ongoing discussions with experiments to understand what information needs to be validated for GLUE 2.
  - ◆ Specific actions are being implemented for ALICE.
  - ◆ Waiting for LHCb migration to GLUE 2 to have more details. So far, they are happy with the existing validation.
  - ◆ No specific requirement from CMS for the time being.
  - ◆ To be understood for ATLAS.

- GOCDB testing instance<sup>?</sup> is now able to filter WLCG services and also services per LHC VOs using the scope option. An option to get T1 and T2 downtimes is under development.
- Ongoing discussions with OIM developers to understand the feasibility of adding more information and implementing similar features as in GOCDB.
- IT-PES has developed the OSG ClassAds to GLUE 2 translator for HTCondor, and together with the MW Officer we are planning the distribution of the rpm through the WLCG repository.
- A TF meeting<sup>?</sup> is scheduled next week where each experiment will present their future interactions with the IS and their plans to migrate to GLUE 2. It will also include a presentation about the GLUE 2 validation status and AGIS.

## IPv6 Validation and Deployment TF

- Deploying an instance of ETF (new implementation of Nagios for SAM) to test the nodes in the IPv6 testbed

Andrea Sciaba clarifies that Nagios is ready to test the IPv6 testbed.

## Middleware Readiness WG

Summary of the 13th WG meeting held on 28 October:

- The xrootd monitoring plugin for dCache v. 2.13.x was installed and tested at PIC together with dCache 2.13.9
- One host of the CERN FTS-3 pilot will be running CentOS7. In this way ATLAS and CMS will be able to test FTS3, via their workflows, on both OS environments.
- DESY has been contacted to arrange ATLAS and CMS tests on their nightly rebuilt dCache instance
- PIC-CERN PhEDEx transfers failing are not yet understood; they are possibly due to a bug in one of the involved MW components (EOS, dCache, FTS-3, Globus, ...) or a misconfiguration somewhere. Experts at CERN are looking into this.
- The MW Readiness App v.3 will move to production real soon now. Volunteer Sites will be called to comment on its functionality.
- There was an ARGUS Collaboration meeting on Oct. 9th. The next one will be on Nov. 6th. The periodical sudden bursts of high load on the CERN ARGUS servers still persist and are not yet explained. A number of other issues from which CMS suffered are now understood. There is one more FTE now working in the ARGUS dev. team. Moving to the upcoming release for CentOS7 may help solve issues, if any, due to historical dependencies: the latest builds use more recent versions of jetty etc.
- Less than usual participants joined this meeting from the Volunteer Sites. Suggested date for the next one is Wednesday 2nd December at 4pm CET. Objections with alternative dates should be sent to [wlcg-ops-coord-wg-middleware@cern](mailto:wlcg-ops-coord-wg-middleware@cern)SPAMNOTNOSPAMPLEASE.ch

Following a question by Christoph Wissing, Andrea Manzi confirms that the dCache xrootd monitoring plugin is available in the WLCG repository.

## Multicore Deployment

## Network and Transfer Metrics WG

- perfSONAR collector, datastore, publisher and dashboard now in production (stable operations)
- perfSONAR 3.5: 205 sonars were updated, ALL sites are encouraged to enable auto-updates for perfSONAR
- Detailed report from the WG presented at GDB [↗](#)
- Meeting held yesterday, encouraging all mesh leaders to participate
- Started discussion on the network outage and at risk announcements from NRENs
- Pilot projects: ATLAS Panda, perfSONAR stream now in ATLAS Network Analytics (<https://twiki.cern.ch/twiki/bin/view/AtlasComputing/ATLASAnalytics>), several KIBANA dashboards available MWT2 [↗](#) FZK2 [↗](#). Jorge and Ilija working on cost matrix using the round-trip time and packet loss in Mathis's formula to infer bandwidth (predictions based on this model will follow).

Alessandro di Girolamo expressed his concerns on exposing site internal information that will be difficult to integrate within the experiment tools that make use of perfSONAR information. For instance, sites comprised of several geographical locations, like GRIF or CERN, expose themselves in a different way. GRIF has one perfSONAR configuration whereas CERN has two. It's not clear this is useful. It is decided to continue the discussion within the WG.

## RFC proxies

- NTR

## Squid Monitoring and HTTP Proxy Discovery TFs

## Action list

Creation date	Description	Responsible	Status	Comments
2015-06-04	Status of fix for Globus library (globus-gssapi-gsi-11.16-1) released in EPEL testing	Andrea Manzi	ONGOING	GGUS:114076 <a href="#">↗</a> is now closed. However, host certificates need to be fixed for any service in WLCG that does not yet work OK with the new algorithm. Otherwise we will get hit early next year when the change finally comes in Globus 6.1.15 GGUS tickets opened for SRM and Myproxy certificates not correct, 6 already closed. OSG and EGI contacted (Maarten alerted the few affected EGI sites as well)
2015-10-01	Follow up on reporting of number of processors with PBS	John Gordon	ONGOING	
2015-10-01	Define procedure to avoid concurrent OUTAGE downtimes at Tier-1 sites	SCOD team	ONGOING	A Google calendar is not yet available, therefore the only way for the moment is to check GOCDB and OIM for downtimes using the links that will be provided for convenience in the minutes of the 3 o'clock operations meeting

## Specific actions for experiments

Creation date	Description	Affected VO	Affected TF	Comments	Deadline	Completion
---------------	-------------	-------------	-------------	----------	----------	------------

## Specific actions for sites

Creation date	Description	Affected VO	Affected TF	Comments	Deadline	Completion
2015-11-05	ATLAS would like to ask sites to provide consistency checks of storage dumps. More information <a href="#">↗</a> and More details <a href="#">↗</a>	ATLAS	-	-	None	-
2015-09-03	T2s are requested to change analysis share from 50% to 25% since ATLAS runs centralised derivation production for analysis	ATLAS	-	-	a.s.a.p.	CLOSED
2015-06-18	CMS asks all T1 and T2 sites to provide Space Monitoring information to complete the picture on space usage at the sites. Please have a look at the general description and also some instructions for site admins.	CMS	-	-	None yet	CLOSED

There is a discussion on the usefulness of the action items defined in the Operations Coordination meetings. Alessandro di Girolamo doesn't find it useful to report about an action that depends on the sites to be completed and not on the experiments. He believes that if the Operations team would be able to follow up on the defined actions it would make sense to define and report about them in the Operations Coordination meetings, otherwise, it's internal work for the experiments that could be tracked internally with less overhead for everyone. Alessandra Forti explains that defining the actions in the meetings helps to spread and advertise the work to be done by the sites. Maria Alandes reminds that this is also included in the sys admin information in the Operations web portal and a link is included in the mail announcing the minutes of the meeting. It is mentioned that few sites attend the meeting and that it's unlikely that people read the minutes up to the end. At the end it is agreed to continue defining action items for sites as a way to summarise and advertise all the open actions in a central way.

## AOB

Michael Ernst raises the issue with the official accounting reports that contain wrong numbers and keep on being produced every month. He would like to know when the new accounting portal will be used, since it shows numbers that are more coherent than the ones in the current production portal. He also stresses the fact that official reports should be correct and not contain wrong information. Alessandra Forti explains that John Gordon reported that the new portal would be online at the end of October but there is no news on this. Maarten Litmaath reminds that this was already raised at the last MB and that Ian Bird was going to follow this up with John Gordon. The Operations Coordination team reminds that the accounting reports are not produced by Operations but they will check with Ian Bird about the status of this.

-- MariaALANDESPRADILLO - 2015-11-03

---

This topic: LCG > WLCGOpsMinutes151105

Topic revision: r20 - 2018-02-28 - MaartenLitmaath



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.



WLCGOpsMinutes151105 < LCG < TWiki

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback