

Table of Contents

WLCG Operations Coordination Minutes, October 1st 2015.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Middleware News.....	1
Tier 0 News.....	2
Tier 1 Feedback.....	2
Tier 2 Feedback.....	3
Experiments Reports.....	3
ALICE.....	3
ATLAS.....	3
CMS.....	3
LHCb.....	4
APEL and multicore accounting.....	5
Ongoing Task Forces and Working Groups.....	5
gLExec Deployment TF.....	5
HTTP Deployment TF.....	5
Information System Evolution.....	5
IPv6 Validation and Deployment TF.....	6
Machine/Job Features.....	6
Middleware Readiness WG.....	6
Multicore Deployment.....	6
Network and Transfer Metrics WG.....	6
RFC proxies.....	7
Squid Monitoring and HTTP Proxy Discovery TFs.....	7
Action list.....	7
Specific actions for experiments.....	8
Specific actions for sites.....	8
AOB.....	8

WLCG Operations Coordination Minutes, October 1st 2015

Highlights

- dCache sites should install the latest fix for SRM solving a vulnerability
- All sites hosting a regional or local site xrootd redirector for CMS should upgrade it at least to version 4.1.1
- CMS DPM sites should consider upgrading dpm-xrootd to version 3.5.5 now (from epel-testing) or after mid October (from epel-stable) to fix a problem affecting AAA
- Tier-1 sites should do their best to avoid scheduling OUTAGE downtimes at the same time as other Tier-1's supporting common LHC VOs. A calendar will be linked in the minutes of the 3 o'clock operations meeting to easily find out if there are already downtimes at a given date
- The multicore accounting for WLCG is now correct for the 99.5% of the CPU time, with the few remaining issues being addressed. Corrected historical accounting data is expected to be available from the production portal by the end of the month
- All LHCb sites will soon be asked to deploy the "machine features" functionality

Agenda

- <https://indico.cern.ch/event/393617/>

Attendance

- local: Maria Dimou (chair), Andrea Sciabà, David Cameron, Marian Babik, Andrea Manzi, Maria Alandes, Maarten Litmaath, Stefan Roiser, Alessandro Di Girolamo, Asa Hsu, Shao-Ting Cheng
- remote: Michael Ernst, Maite Barroso Lopez, Hung-Te Lee, Renaud Vernet, Thomas Hartmann, Di Qing, Antonio Maria Perez Calero Yzquierdo, Stuart Pullinger (APEL), Christoph Wissing, Catherine Biscarat, Daniele Cesini, Pepe Flix, David Mason, Massimo Sgaravatto, Vincenzo Spinoso, Rob Quick, Alessandra Doria, Alessandra Forti, Gareth Smith, John Gordon, Jeremy Coles

Operations News

The date of the next WLCG Operations Coordination meeting was moved from 15/10 (HEPiX @ BNL & GDB week) to 22/10.

Middleware News

- Useful Links:
 - ◆ Baseline Versions
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines:
 - ◆ dCache 2.10.41/2.12.21/2.13.9 released to fix a vulnerability in SRM
 - ◆ XRootD 4.1.1 at least in regional and local site redirectors has been requested by both ATLAS and CMS for FAX and AAA mainly given the increasing number of clients using IPv6. AAA is planning to put all site with 3.3.x version to a transitional federation and give limited support to them. FAX recommends 4.2.3 for redirectors

- ◆ FTS 3.3.2. Configuration changes and fixes to reduce the django memory footprint. The memory problem affected RAL 2 weeks ago due to the high number of transfers submitted.
- Issues:
 - ◆ CMS AAA issues with DPM. DPM-Xrootd wrongly sends the XrdCmsClient Remove notices on ENOENT, this condition can be triggered if a remote client includes a "tried" list that indicates the site has been tried before. Having the cmsd then send a removed notice can cause problems, as reported by AAA (cms) xrootd federation. A new version of dpm-xrootd (3.5.5.) with the fix has been pushed to epel-testing and verified for SL6 (there is a problem internal to Fedora system which is locking the SL5 update though), SL6 sites can either upgrade from epel-testing now or wait from the push to epel stable (mid october)
 - ◆ the latest version of HTCondor (8.4.0) cannot be installed on a WN installed from UMD3 due to incompatible dependency, (HTCondor depends now on condor-classads, while WN depends on glite-lb-common which depends on classads). Discussion is ongoing on how to solve the issue at GGUS:116274 [↗](#)
- T0 and T1 services
 - ◆ BNL
 - ◇ dCache upgraded to 2.10.39 (gridftp nodes) and 2.10.41 (srm)
 - ◇ XRootD upgrades
 - Reverse Proxy 4.1.1
 - Forward Proxy Dev Version 4.2.3
 - North American redirectors 4.2.1
 - ◇ FTS3 upgraded to 3.3.1
 - ◆ CERN
 - ◇ FTS3 upgraded to 3.3.2
 - ◆ CNAF
 - ◇ FTS3 upgraded to 3.3.1
 - ◆ IN2P3
 - ◇ dCache upgraded to 2.10.41
 - ◆ RAL
 - ◇ FTS3 upgraded to 3.3.2
 - ◆ RRC-KI-T1
 - ◇ dCache upgraded to 2.10.41

Tier 0 News

LSF intervention: We are planning to migrate the CERN LSF masters to a new version of LSF, from 7.06 to 9.1.3. The main motivation for this migration is the end of life of the current production version; there will be no changes to the end-user functionality. The intervention is planned to start sometime after the week of 12/10/2015, and it will consist of two service interruptions of about 20-30min each during which it will not be possible to submit new jobs or query the status of running jobs. Running jobs will continue to run, pending jobs will not be dispatched during the two service interruption windows.

The exact date of the intervention will be announced in due time.

Please, let us know as soon as possible if this causes any problem to your planned activities on that time frame.

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- normal to high activity
 - ◆ Reco, analysis, MC
 - ◆ Final preparations for Quark Matter 2015[☞](#), Sep 27 - Oct 3
- CERN
 - ◆ Since ~3 weeks the reco jobs have been downloading raw data files from CASTOR instead of streaming them.
 - ◇ Looks good.
 - ◆ Further meetings with CASTOR and EOS experts to discuss short- and longer-term ideas to stabilize usage of CASTOR.
 - ◇ Various scenarios are available to help ensure a successful heavy-ion data taking period.
 - ◇ T0ALICE is being enlarged with more disk servers (see below).
 - ◆ Timeouts and failures hampered data transfers from P2 to CASTOR starting Mon Sep 21 afternoon:
 - ◇ Initial investigations got sidetracked.
 - ◇ Data taking had to be stopped Tue morning when all P2 disks were full.
 - The beam happened to get dumped half an hour later.
 - ◇ The trouble was due to new disk servers having an incorrect setup.
 - ◇ The CASTOR team fixed that within ~1 hour.
 - ◆ We thank the CASTOR team for their good support in these matters!
- A successful workshop on network evolution in Asia[☞](#) was held at KISTI Sep 22-24.
 - ◆ LHCONE peering was agreed between TEIN, KREONET2 and ASGC.
 - ◆ We thank the KISTI team, the network experts, and the site admins for making this happen!

ATLAS

- Very high activity on Grid and a large increase in data taking from the detector
- FTS problems affected data transfer and jobs in the last 2 weeks
- Some workflow changes being rolled out:
 - ◆ Leaving output datasets where they are produced instead of consolidating at one site
 - ◆ At T2s, merging the job in/out buffer spacetoken (PRODDISK) with persistent data spacetoken (DATADISK) - this allows T2s to be final destination of data (currently only T1s)
 - ◆ Removing further the "cloud" boundaries (T1 and associated T2s)
 - ◆ Intelligent job retries, not retrying permanent errors
 - ◆ Transfer source selection and queuing in Rucio rather than FTS

CMS

- Main operational issues are "CMS internal"
 - ◆ Run into some large backlog of file registrations for CMS Tier-0 application
 - ◆ High load on CMS DBS (Data Bookkeeping Systems)
 - ◆ Issues partly related
- Multi-Core accounting
 - ◆ Issue for ARC-CEs with Torque/PBS schedule GGUS:114382[☞](#)
- Request to make DPM 1.8.10 minimum in baseline
 - ◆ Earlier versions have issues to work in CMS xrootd Federation
- Ticketing campaign for proper AAA (CMS xrootd Federation) configuration at sites

- ◆ Proper publishing of monitoring
- ◆ Proper configuration to allow scale test running
- Still ongoing investigations about SAM test scheduling at Cern
 - ◆ Affects tests submitted with role 'pilot', GGUS:116468 [↗](#) and GGUS:116069 [↗](#)
 - ◆ Long lasting empty pilots from CMS under investigation

Andrea M. reports that some French sites are already running DPM with the fix installed without any issue. It is agreed anyway that the change in baseline version will happen only when the new version is in production and verified to be working fine for some time, which means that it will not happen before three weeks. This is not a problem because many sites will upgrade before then.

LHCb

- Operations
 - ◆ Run 2 data processing in full swing and progressing well. Currently data processing activities are limited to T0/1 sites but can be extended to T2 sites if needed (mesh processing).
 - ◆ Very high activity, LHCb using all available distributed compute resources available to the VO
- Issues
 - ◆ File export from pit occasionally hangs with SRM_BUSY errors in Castor leaving files in unknown state and need to be cleaned up by hand by data management before retrying a transfer. In contact with Castor team on finding a solution.
 - ◆ This week 3 T1 sites happened to be not available for LHCb data processing on Tuesday. All announcements were done in GOCDB properly and/or via mail announcements and site contacts. Following GOC downtimes only 2 sites would have been in DT at the same time, the third site became not available because it started draining 24 hours before the DT (which was also announced). Requesting to WLCG as a coordinating body of distributed compute resources - after a similar case earlier this year - to coordinate downtimes for LHCb supporting T1 sites to optimally have 1 site in scheduled DT at a time, two are a burden. Excepted are of course urgent scheduled downtimes.
 - ◇ Reminder, if sites stop submission because of draining before a downtime please close the CE instead of disabling submission to batch, not resulting in aborted pilots in the monitoring.
 - ◆ Over last week-end (3 consecutive days) a decline of CERN/LSF running jobs was observed. As of Tuesday the normal resetting of the fair share early in the morning and therefore increased number of running jobs can be seen again.
 - ◆ Because of an issue with the LHCbDIRAC CVMFS installation the necessary installation had to be done via download of the package from a remote server. It was observed that the network link to RRCKI had problems with supporting this installation type. In the meantime CVMFS installation is fine and therefore the issue is not seen anymore.
 - ◆ "Slow CERN/LSF workernodes" has been looked into by the PES team and it seems the cause of the problem is understood and will be fixed, many thanks for the efforts.
- Developments
 - ◆ Working on a HTCondor CE submission. To be tested with the CERN HTCondor infrastructure.

It is agreed that all Tier-1 sites declared their downtimes following the correct procedures, however in the future concurrent downtimes should be avoided.

Andrea S. reports that the calendars for the downtimes will appear in the minutes for the 3 o'clock operations meetings from now on. The functionality is considered a useful tool.

The SCOD will check if there are clashes and do his/her best to prevent them, removing them by asking sites to reschedule or in the worst case to point them out very clearly.

Concerning the HTCondor plugin, Maarten reports that also ALICE are working on one, already tested with a "hello-world" job.

About the "slow WN" problem, Maite clarifies that it was due to some optimisation of the memory management of the hypervisor, which did not work as expected. The fix has been applied to just a few nodes and they are waiting for LHCb to confirm that it works fine.

APEL and multicore accounting

John G. reports the status of multicore accounting. Currently 99.5% of the WLCG CPU time comes from sites correctly reporting the number of cores. Most of the remaining 0.5% is due to DESY, which reports zero cores, due to a problem in the ARC-PBS combination, which is tracked in GGUS:114382[\[7\]](#), while CREAM CEs at DESY and other PBS sites report Processors=1 for multicore jobs (now fixed at DESY). According to the plans, the dev portal will contain all historical APEL data with cores reported only from when sites started publishing them; when this is done, T1 and T2 reports will use this data. The portal is undergoing a major rewrite that will be released in April.

Antonio reports that from the CMS point of view the situation is clear, only two sites are not reporting cores, DESY and Kharkov.

Maarten and Alessandro point out that for WLCG the dev portal should be "promoted" to production, and Pepe adds that the wrong reports generated by the production portal are creating a lot of confusion for the management and the referees. John answers that by the end of October all the T1/T2 reports will show all the corrected historical data.

Pepe remarks that the official WLCG accounting reports were generated using the wrong data from the production portal and suggests to regenerate them. John points out that the T1 data needs to be completely corrected before doing that; he proposes to ask the MB if the reports need to be regenerated.

Ongoing Task Forces and Working Groups

gLExec Deployment TF

- NTR

HTTP Deployment TF

Information System Evolution

- There was a TF meeting[\[8\]](#) where the following presentations were made:
 - ◆ Follow up on MB&GDB presentations:
 - ◇ it was agreed to investigate the possibility of using OIM/GOCDB as service registries, extending the information they currently provide to meet use cases for static/mutable information; and to query the resource BDII/OSG collectors for dynamic information.
 - ◇ it was also agreed to consider the implementation of a WLCG profile to target the validation of information on WLCG use cases. Discussions on going with EGI to consider the integration of glue-validator at resource BDII level.
 - ◆ OSG presented their plans to move to ClassAds and OSG collectors to provide information about their resources, for the time being for HTCondorCE. A translator from ClassAds to GLUE 2 is developed by OSG and CERN IT-PES. MW Officer in contact with developer at

CERN to understand how this translator could be distributed to all sites.

- ◆ EGI presented their plans where the current information system is going to be used with the idea of moving to GLUE 2 and deprecating GLUE 1, as long as WLCG doesn't depend on it. It was agreed to plan for a transition in WLCG so that GLUE 2 information is consumed and we stop relying on GLUE 1.
- ◆ NDGF presented the way in which they currently publish information, supporting both nordugrid and GLUE 2 schemas. They would prefer if we could move to GLUE 2 to make things simpler.
- ◆ GOCDB developer made a presentation of technical details and features available in GOCDB that would allow us to move in the service registry direction.
- Ongoing discussions in the mailing list to resurrect ginfo as the GLUE 2 client tool to query the information system.

IPv6 Validation and Deployment TF

Machine/Job Features

Stefan reports that after the GDB it has become clear that the MJF functionality is interesting only for LHCb, at the moment, so a campaign will be launched to ask only the LHCb sites to deploy the machine features part. LHCb will collect experience on its use cases and report back to WLCG. It is agreed that in future meetings we can remove the MJF TF from the agenda and have reports only when the need arises.

Middleware Readiness WG

- A puppet-pakiti module configuring the pakiti-client cron with the parameters needed for WLCG MW Readiness is available at <https://gitlab.cern.ch/wlcg-mw-readiness/puppet-pakiti>
- StoRM testing for ATLAS in process at INFN_T1. Failing jobs under investigation. JIRA:MWREADY-61
- The dCache 2.13.x xrootd monitoring plugin has been prepared by the developer Ilija Vukotic, tested by PIC and pushed to the WLCG repo.
- DPM 1.8.10 verification for CMS completed at GRIF. JIRA:MWREADY-83
- DPM 1.8.10 verification on CentOS7 started at GLASGOW. Now arranging a test set-up for the ATLAS workflow JIRA:MWREADY-82
- EOS 0.3.129-aquamarine verification pending for CMS at CERN JIRA:MWREADY-81
- Next meeting reminder **October 28th at 4pm CET**. Agenda http://indico.cern.ch/e/MW-Readiness_13

Multicore Deployment

Network and Transfer Metrics WG

- Meeting held yesterday, <https://indico.cern.ch/event/400643/>
- Publishing of the perfSONAR results using OSG production service planned for 13th of October (OSG production date)
- OSG dashboard (psmad.grid.iu.edu) will go production on the same date, already showing more recent results than maddash.aglt2.org, one issue to be fixed is to correctly show tests done in one-direction only

- WLCG-wide meshes campaign finalized with 94 sonars in latency testing, 115 sonars in traceroutes and 104 in throughput.
 - ◆ Sonars that were not included in the WLCG-wide meshes were reported to the mesh leaders and will be followed up (currently they reside in the global meshes, once issues are fixed they'll be moved to WLCG meshes)
 - ◆ Started re-creating project meshes, Belle II and Dual-stack (IPv4/IPv6 bandwidth), plans for other meshes to be discussed
- Once infrastructure is in production, we plan to focus on the integration projects, there are ongoing pilot projects for ATLAS and LHCb
- There is also interest in perfSONAR in the IT Analytics WG as well as from the network community Asia Tier Centre Forum (<https://indico.cern.ch/event/395656/>)
- perfSONAR 3.5 was released on Monday 28th Sept, 162 sonars were auto-updated, 68 still on 3.4, all sites are encouraged to enable auto-updates for perfSONAR
- Next WG meetings will be on 4th of Nov and 2nd of Dec

RFC proxies

- NTR

Squid Monitoring and HTTP Proxy Discovery TFs

- Nothing to report

Action list

Creation date	Description	Responsible	Status	Comments
2015-09-03	Status of multi-core accounting	John Gordon	CLOSED	A presentation about the plans to provide multicore accounting data in the Accounting portal should be presented at the next Ops Coord meeting on October 1st https://indico.cern.ch/event/393617/ since this is a long standing issue
2015-06-04	Status of fix for Globus library (globus-gssapi-gsi-11.16-1) released in EPEL testing	Andrea Manzi	ONGOING	GGUS:114076 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. A broadcast message has been sent by EGI. Now the team will start working on the monitoring script that will show the sites that haven't changed and open GGUS tickets to remind them. JIRA:MWREADY-86 is now closed
2015-10-01	Follow up on reporting of number of processors with PBS	John Gordon	CREATED	
2015-10-01	Define procedure to avoid concurrent OUTAGE downtimes at Tier-1 sites	SCOD team	CREATED	

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
2015-09-03	T2s are requested to change analysis share from 50% to 25% since ATLAS runs centralised derivation production for analysis	ATLAS	-	-	a.s.a.p.	ONGOING
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	~10 T2 sites missing, Ticket open

AOB

-- MariaDimou - 2015-09-14

This topic: LCG > WLCGOpsMinutes151001

Topic revision: r52 - 2018-02-28 - MaartenLitmaath



Copyright &© 2008-2024 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