

# Table of Contents

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

# WLCG Operations Coordination Minutes - October 2nd, 2014

## Agenda

- <https://indico.cern.ch/event/341734/>

## Attendance

- local: Nicolo Magini (secretary, CMS), Maria Alandes, Marian Babik, Felix Lee (ASGC), Romain Wartel (security officer), Andrea Manzi (MW officer), Andrea Sciaba, Simone Campana (ATLAS), Stefan Roiser (LHCb), Alberto AImar, Marcin Blaszczyk, Manuel Guijarro (Tier-0), Maarten Litmaath (ALICE), Maria Dimou
- remote: Alessandra Forti (chair), Andrej Filipcic (ATLAS), Antonio Maria Perez Calero, Burt Holzman (FNAL), Catherine Biscarat, Di Qing, Ewan Mac Mahon, Jeremy Coles (GridPP), Massimo Sgaravatto, Michael Ernst (BNL), Peter Solagna (EGI), Renaud Vernet, Rob Quick (OSG), Sang Un Ahn, Shawn Mckee, Ulf Tigerstedt, Ron Trompert (NL-T1), Pepe Flix (PIC)

## Operations News

- HEP\_OSlibs-7.0.0-0.el7.cern.x86\_64.rpm for CentOS7 has been released. The metarpm is based on the previous SL6 one with packages not present anymore in CERN CentOS7 removed. More information:
  - ◆ <https://twiki.cern.ch/twiki/bin/view/LCG/CentOS7DependencyRPM>
- CHEP2015 deadline for abstracts in on the 15th of October 2014.
  - ◆ <https://indico.cern.ch/event/304944/call-for-abstracts/>
- The Shellshock vulnerability implies a set of security bugs affecting the bash shell. This was disclosed on 24th September and the WLCG Security experts has been evaluating since then the impact of these vulnerability in the WLCG infrastructure. Some perfSONAR nodes have been found to be compromised. The security teams, as well as WLCG Operations, **highly recommend that all sites terminate their perfSONAR instances** as a precautionary measure, until the attacks are contained, unless you have patched the Bash packages on the perfSONAR instance(s) by Friday 26 Sep and can actively confirm by checking network logs that NO IRC traffic was emitted from your hosts. Investigations of other services are ongoing. WEB servers are particularly targeted.
- Rob Quick reports that one OSG BDII was compromised via Shellshock, the incident was reported to the OSG security team. The OSG BDII was compromised through apache, which was running to provide a human-readable interface: since this is not normally the case for BDII, the issue doesn't affect BDII in general. The server was patched and will be reinstalled from scratch soon. Romain reminds that any web service is at risk and should be upgraded now.
- Peter Solagna presents a proposal for time limit normalisation in VOcards, see the agenda  for details.
- Maarten asks if a new field for HEPSPEC\*time limit could be added instead of replacing the existing time limit, this would ease the transition for ALICE; the old field could be phased out later. Peter is worried that sites won't know which one to use; Stefan and Andrea Sciaba suggest to allow selecting only one of the two fields. Discussion to be continued offline.

## Middleware News

- Useful Links:
  - ◆ Baseline Versions
  - ◆ MW Issues
  - ◆ Storage Deployment
  
- Baselines:
  - ◆ News from EGI URT meeting on monday
    - ◇ New version of the UI and WN are going to be hopefully included in the next version of UMD foreseen for end of October.
    - ◇ dCache 2.2.x decommissioned deadline is 31-10-2014. Baseline for the 2.6.x series is 2.6.34 which fixes issues with Brazilian CAs certs.
    - ◇ Globus 6 is in epel-testing and PTs are invited to test compatibility. We are aware of FTS and DPM being tested, at the moment not blocking issues have been discovered and therefore soon Globus 6 will go to stable. ( the exact date will be discussed during the next URT)
  
- MW Issues:
  - ◆ xroot package deployed with ROOT 6 breaks access to dCache storage, affecting LHCb. The problem is both client and service side, A fix for dCache has been developed but not yet released, at the moment there will be a workaround fix in ROOT.
  - ◆ installation of several grid products is broken. CREAM, WMS, L&B, UI, WN cannot be installed at the moment cause the classads package ( dependency for all of them ) was declared orphan in EPEL, and retired from the EPEL repository. The package is going to be included at the moment into UMD and EMI third-party repo, waiting for a maintainer. CESNET should take care of it but they are not happy with this extra effort.
  
- About the issue of xroot with ROOT6, Stefan comments that the client-side workaround is available; still waiting for server-side fix by dCache.
- Maria Alandes comments that the classADS package was owned by Steve Traylen in EPEL. Andrea Manzi and Manuel Guijarro comment that he was admin of the package (coming from Condor) in EPEL, but not maintainer. Alessandra suggests that since the affected services we care about are CREAM, WN and the UI the CREAM team could take over the package. Andrea Manzi will follow the discussion on this issue at the next EGI URT meeting.
  
- T0 and T1 services
  - ◆ IN2P3
    - ◇ dCache upgrade to 2.6.34
  - ◆ NL-T1
    - ◇ dCache upgrade to 2.6.34
  - ◆ KIT
    - ◇ Update of dCache for CMS and LHCb to 2.6.34
    - ◇ Update xrootd configuration for FAX and AAA to respect EU privacy policy Thursday 08:00 - 08:30 UTC.
    - ◇ Update for LHCb dCache to next version that fixes issues with ROOT6 not scheduled yet (new dCache release required first).
  - ◆ JINR-T1
    - ◇ one dCache instance upgrade to 2.6.34
    - ◇ one dCache instance running 2.2.27 to be upgraded to 2.6 or 2.10 early november
  - ◆ BNL
    - ◇ FTS upgrade to 3.2.27

## Oracle Deployment

- IT-DB new hardware installations in: CERN computer centre and Wigner.
- Timeline: testing in October, production move - by the end of 2014. Schedule will be updated accordingly.
- Following table includes only those DB services that concern WLCG

Database	Comment	Destination	Upgrade plans, dates
ATONR	Data Guard for Atlas Online	Wigner	
ATLR	Data Guard for Atlas Offline	Wigner	
ADCR	Data Guard for ADC	Wigner	
CMSR	Data Guard for CMS offline	Wigner	
LHCBR	Data Guard for LHCB Offline	Wigner	
LCGR	Data Guard for WLCG	Wigner	
CMSONR	Data Guard for CMS Online	Wigner	
LHCBONR	Data Guard for LHCB Online	Wigner	
CASTORNS	Data Guard for Castor Nameserver	Wigner	
ATONR	Active Data Guard for Atlas Online	CERN CC	
ALIONR	Active Data Guard for Alice Online	CERN CC	
ADCR	Active Data Guard for ADC	CERN CC	
LHCBONR	Primary Database for LHCB Online	CERN CC	

- Upgrade Plans Table from Q1&Q2 2014
- Marcin reminds that the Data Guard DBs are only used for disaster recovery, so the intervention will not affect users. The Active Data Guard DBs are accessed by users in read-only mode.

## Tier 0 News

- **AFS UI:** Waiting feedback from the experiments (see action list)
- **WMS** service was decommissioned on October 1st
- **plus5/batch5:** user feedback fully analyzed, no major showstoppers. Some alternatives discussed with different user groups. Lxplus5 will be stopped in October, exact date being discussed.
- Next job efficiency meeting on October 10th,  
<https://twiki.cern.ch/twiki/bin/view/PESgroup/MeetingHeld10thOct2014>
- Simone comments that ATLAS is not happy with the retirement on October 31st of machines in SafeHost running ATLAS services which cannot/will not be migrated to AI by that date. Manuel answers that the CERN hosting in SafeHost will end in spring; CERN wants to start the retirement now but it's not a problem to keep some machines there until spring, send a list of affected machines. Other VOs also sent similar concerns.
- Stefan asks if the AFS UI access logs can also include user names, Manuel answers that they only have hostnames. Maarten suggests to move the UI to a different space instead of deleting it. Stefan, Maarten and Nicolo remind that some VO services still use the clients from the AFS UI, not just the CRL distribution.

## Tier 1 Feedback

## Tier 2 Feedback

## Experiments Reports

### ALICE

- CERN: investigation of job failure rates and inefficiencies
  - ◆ the batch team adjusted the parameters of the ALICE queue - **thanks!**
  - ◆ the effects should become visible over the coming days
  - ◆ the pilot ("job agent") logic has been checked and a potential improvement is being looked into
- CERN: HLT farm [↗](#) running as an ALICE site since Sep 24
  - ◆ being ramped up over the coming months

### ATLAS

- DC14
  - ◆ digi+reco Run2 in 8-core mode will finish in about a week from now
  - ◆ some more simulation samples launched
  - ◆ further AOD2AOD on reprocessed data will be launched in one week (0.5PB)
- Multicore recommendations for 8-core reconstruction
  - ◆ Allocate 16GB physical memory per job
  - ◆ if limiting memory per process: 3GB RSS and/or 5GB VMEM
  - ◆ For cgroup-enabled sites: total RSS of the job should be 16GB
- Deployment of ATLAS MCORE queues
  - ◆ more than 70k cores were used this week for multicore jobs exclusively
  - ◆ after fixing initial memory issues, the digi+reco processing is progressing very fast, up to 50M events/daily - thanks to the sites for fast action
  - ◆ to all the sites, please continue the deployment of multicore queues
  - ◆ serial production tasks in the future will be limited
- SAM3 tests
  - ◆ all ARC-CE sites fixed the configuration and the ATLAS\_CRITICAL tests are effective since 1.10.
- Rucio and Prodsys-2 commissioning ongoing, still no fixed date for deployment

### CMS

- Processing overview:
  - ◆ Not much work in the system
  - ◆ Preparations for new campaign (PHYS14) ongoing
- Scale testing of HTCondor and GlideinWMS by OSG colleagues
  - ◆ Launch many pilots on one acquired job slot to reach high scales
  - ◆ Caused some trouble at sites
    - ◇ Firewalls not able to handle that many connections
    - ◇ Maximum number of NAT connections exhausted
    - ◇ Report problems to CMS (via ticket or HyperNews) - we negotiate with the testers
- Problems with Dashboard reporting
  - ◆ Dashboard team presently working on a job monitoring collector with UDP
- AFS-UI at CERN
  - ◆ Analysis of AFS access logs
    - ◇ ~45 individual analysis users - usage likely decrease after closing of lxplus5
    - ◇ ~5 users from central production
  - ◆ Extending the UI availability beyond Oct 31st is still preferred
    - ◇ Perhaps even needed - migration path of some services still to be understood
- Reminders for sites

- ◆ Participate in space monitoring (compare last meeting)
  - ◆ Update xrootd fallback configuration
    - ◇ Opened tickets to various sites - quite some took action - thanks!
  - ◆ Add "Phedex Node Name" to site configuration
- 
- Nicolo' asks about the procedure to enable the CMS services to get proxies from the new VOMS servers for testing. Maarten suggests to send a mail to himself and the service managers (Alberto Peon, Steve Traylen)

## LHCb

- Access to dCache storage sites is broken when accessed by ROOT6/xrootd, i.e. a negotiation for a vector read fails and subsequently ROOT crashes. A fix is proposed by the dCache team. On the ROOT side an intermediate stop gap solution is possible to be deployed until the dCache fix is out and deployed.
- AFS UI references have been checked and eventually cleaned from all LHCb distributed computing clients and tools. The retirement of the UI should be possible for LHCb.
- It was discovered that WLCG reports for CERN not only contain statistics for worker nodes but all resources used by the experiment (including VOBOXes, build nodes, etc.). This makes it hard to compare e.g. job efficiencies to other sites and LHCb proposed to only publish worker node figures (e.g. WLCG report, page 48 [↗](#))
- A new stripping campaign is currently being prepared by LHCb. This campaign will produce a "legacy dataset" for 2010-12 data. The plan is to also do partly reconstruction and include tagging information which will result in more work to be executed. The net result for the sites is that the staging will most likely not be the bottleneck for this operation.
- New VOMS servers are currently being tested in certification by LHCbDIRAC with full workflows.

## Ongoing Task Forces and Working Groups

### gLExec Deployment TF

- NTR

### Machine/Job Features

- NTR

### Middleware Readiness WG

The 6th meeting of the WG took place yesterday, as planned. Please follow the presentations of the MW Officer Andrea M. and the MW Package Reporter developer L.Cons from the agenda [HERE](#) [↗](#) to see the products in the pipeline for Readiness verification and the the one of the developer L.Cons scenarios for the Collector/Reporter under evaluation. Actions were completed. Next meeting on Wed Nov 19th at 4pm CET. Please do note the date/time!

### Multicore Deployment

- ATLAS: see ATLAS report

### SHA-2 Migration TF

- introduction of the new VOMS servers

- ◆ node firewalls are being opened selectively per experiment
  - ◇ ALICE OK
  - ◇ LHCb looking good
  - ◇ CMS in progress
    - FNAL FTS-3 config needed fixing
  - ◇ ATLAS in progress
    - BNL FTS-3 config needed fixing
- ◆ EGI will soon conclude their campaign to get all EGI sites to recognize the new servers for the Ops VO
  - ◇ the port for Ops will be opened at that time
- ◆ our special routing rules have been extended until Tue Nov 18 (sic)
  - ◇ those rules allow remote sites to get "Connection refused" instead of timeouts
  - ◇ by that time we still have 1 week to fix unwanted behavior
  - ◇ we should have things in good shape long beforehand...

## WMS Decommissioning TF

- Marian comments that with the deployment of the Condor SAM probes, nothing is using WMS anymore. Alessandra comments that the SAM probe transition went well for ATLAS.
- Manuel reports that the machines were stopped, and will be destroyed in 2 weeks.
- Agreed to CLOSE the task force.

## IPv6 Validation and Deployment TF

- IPv6 tests done for LHCbDIRAC, network configuration of LHCbDIRAC and authentication of services/agents across different machines is working

## Squid Monitoring and HTTP Proxy Discovery TFs

- No progress to report again this meeting

## Network and Transfer Metrics WG

- Details on the shell shock vulnerabilities and its impact on perfSONAR available at <https://twiki.cern.ch/twiki/bin/view/LCG/ShellShockperfSONAR>
- We recommend ALL sites that didn't patch bash before Friday Sep 26 to terminate their instances and wait until perfSONAR 3.4 is released
- perfSONAR 3.4 to be released on Mon Oct 6, WLCG and EGI broadcasts will be sent with the installation instructions
- perfSONAR operations meeting this Friday (Oct 3 at 3PM), agenda at <https://indico.cern.ch/event/342995/>
- Discussing how to mitigate the risk of exposure of PerfSONAR using iptables. Anyway the reappearance of similar issues cannot be ruled out completely since PerfSONAR needs to be open to traffic.

## Action list

1. CLOSED on the WLCG middleware officer: to take the steps needed to enable the CVMFS UI distribution as a viable replacement for the AFS UI.
  - ◆ The CVMFS [grid.cern.ch](https://grid.cern.ch/emi-ui-3.7.3) contains the emi-ui-3.7.3 SL6 (path `/cvmfs/grid.cern.ch/emi-ui-3.7.3-1_sv6v1`) and provides as well CA certs, crls and voms lsc files. Given the new UI release we can also plan to upload the UI v3.10.0.

- ◆ TODO: clarify the responsibilities (including ticketing etc.) for the maintenance of the CVMFS UI, in particular running fetch-crl
- ◆ UPDATE: Steve said that the grid.cern.ch CVMFS server maintenance is under PES responsibility, so also the fetch-crl update process. In case of issues the Configuration Management SE should be addressed.
  - ◇ Manuel comments that the CVMFS SE can also be used.
- 2. ONGOING on the experiments: check the usage statistics of the AFS UI, report on the use cases at the next meeting.
- 3. ONGOING on Andrea S.: to understand with EGI if it is possible to bypass the validation of VO card changes in the case of the LHC VOs. Peter Solagna asked the developers to implement the change, Andrea will check status for next meeting.
- 4. ONGOING on the WLCG monitoring team: evaluate whether SAM works fine with HTCondor CE. Report about showstoppers. Status: the SAM team made a proposal on the steps to taken to enable SAM. ATLAS is following up to make sure that the new CEs are correctly visible in AGIS, while for the CMS VO feed they will be taken directly from OIM. The plan is at first to test HTCondor-CEs in preproduction and later switch to production. It is not foreseen to monitor at the same time GT5 and HTCondor endpoints on the same host.
  - ◆ No showstopper for SAM. Need to discover topology; publishing queues in BDII not necessary for SAM probes since Condor can choose the queue based on the proxy.
- 5. CLOSED on Alessandro DG: find out from OSG about plans for publication of HTCondor CE in information system, and report findings to WLCG Ops. To be followed up with Michael Ernst and Brian Bockelman.
  - ◆ Michael Ernst reports that OSG is enabling the info collector, working with ATLAS to get the requirements. Rob Quick comments that info will be published only in Glue1 schema.
  - ◆ Maarten asks if CMS was involved, Rob answers that Brian is in the collector project and CMS expressed no concern so far.

## AOB

- Next meeting on October 16th

-- NicoloMagini - 19 Sep 2014

---

This topic: LCG > WLCGOpsMinutes141002

Topic revision: r30 - 2014-10-06 - JosepFlix



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