

# Table of Contents

<b>WLCG Operations Coordination Minutes - 30 May 2013.....</b>	<b>1</b>
Agenda.....	1
Attendance.....	1
News.....	1
Middleware news and baseline versions (N. Magini).....	1
Tier-1 Grid services (N. Magini).....	2
Storage deployment.....	2
FTS deployment.....	3
LFC deployment.....	3
Other site news.....	3
Data management provider news.....	4
Machine/job features (S. Roiser).....	4
Experiment operations review and plans.....	4
ALICE (M. Litmaath).....	4
ATLAS (A. Di Girolamo, I. Ueda).....	4
CMS (C. Wissing).....	5
LHCb (J. Closier, S. Roiser).....	5
Task Force reports.....	5
SL6 (A. Forti).....	5
gLExec (M. Litmaath).....	6
SHA-2 migration (M. Litmaath).....	6
FTS-3 (M. Salichos, A. Di Girolamo).....	7
perfSONAR (S. Campana).....	7
WLCG monitoring support unit (M. Babik).....	7
AOB.....	7
Action list.....	8
Chat room comments.....	8

# WLCG Operations Coordination Minutes - 30 May 2013

## Agenda

<https://indico.cern.ch/conferenceDisplay.py?confId=254445>

## Attendance

- Local: Maria Girone (chair), Andrea Sciabà (secretary), Marian Babik, Oliver Gutsche, Christoph Wissing, Michail Salichos, Oliver Keeble, Alessandro Di Girolamo, Jan Iven, Felix Lee, Ikuo Ueda, Simone Campana, Nicolò Magini, Andrea Valassi, Ian Fisk, Pablo Saiz, Maite Barroso, Maria Dimou, Maarten Litmaath
- Remote: Stefan Roiser, Joel Closier, Alessandro Cavalli, Michel Jouvin, Di Qing, Stephen Burke, Javier Sánchez, Luís Emiro Linares García, Jeremy Coles, Peter Solagna, Alessandra Forti, Rob Quick, Josep Flix, Isidro González Caballero, Peter Gronbech

## News

Pepe: Isidro and Javier joined the working group as representatives for the Spanish sites.

Michel asks if at the next GDB it is possible to discuss the WebDAV deployment needed by ATLAS for RUCIO. Ueda thinks that even if ATLAS is in close contact with its sites via internal communication channels, the GDB and the WLCG ops coordination meeting can be useful to make this need even more official. The only problem is that next GDB overlaps with the ATLAS offline and software week.

## Middleware news and baseline versions (N. Magini)

<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGBaselineVersions>

Highlights:

- the top BDII version was updated to fix some problems in GLUE-2 support and upgrade is recommended
- the CVMFS version was updated
- There is a security fix for the WMS

Michel asks what is the position of the WLCG operations coordination WG concerning upgrading to EMI-3. Maarten explains that the current recommendation is to upgrade for specific services, the list of which will gradually increase, but sites do not have to upgrade everything to EMI-3. A motivation for a full upgrade could be for example to reduce the support effort for EMI-2. Currently the only strong driving forces to upgrade to EMI-3 are SHA-2 and GLUE-2. A strong push now is not recommended until we have more solid evidence that there are no issues from the experience of sites running EMI-3 services. Alessandra reports that in any case a number of sites have already moved various services to EMI-3. Maarten hopes that sites will proactively upgrade without the need for formal coordination.

Special care must be taken for the worker nodes: we should be sure that nothing breaks from EMI-2 to EMI-3 and a more formal testing from the experiments is advisable.

The WLCG operations report in the next GDB will include this topic.

## Tier-1 Grid services (N. Magini)

### Storage deployment

Site	Status	Recent changes	Planned changes
CERN	CASTOR: v2.1.13-9 and SRM-2.11 for all instances EOS: ALICE (EOS 0.2.29 / xrootd 3.2.7) ATLAS (EOS 0.2.28 / xrootd 3.2.7 / BeStMan2-2.2.2) CMS (EOS 0.2.29 / xrootd 3.2.7 / BeStMan2-2.2.2) LHCb (EOS 0.2.29 / xrootd 3.2.7 / BeStMan2-2.2.2)		
ASGC	CASTOR 2.1.13-9 CASTOR SRM 2.11-2 DPM 1.8.6-1 xrootd 3.2.7-1		
BNL	dCache 1.9.12.10 (Chimera, Postgres 9 w/ hot backup) http (aria2c) and xrootd/Scalla on each pool	Upgrade to dCache 2.2-10 May 28, 2013	None
CNAF	StoRM 1.8.1 (Atlas, CMS, LHCb)		Tests still ongoing, upgrade to StoRM 1.12.0 planned for mid June
FNAL	dCache 1.9.5-23 (PNFS, postgres 8 with backup, distributed SRM) httpd=2.2.3 Scalla xrootd 2.9.7/3.2.7-2.osg Oracle Lustre 1.8.6 EOS 0.2.30-1/xrootd 3.2.7-2.osg with Bestman 2.2.2.0.10		Upgrade to dCache 2 + Chimera this summer
IN2P3	dCache 1.9.12-16 (Chimera) on SL6 core servers and 1.9.12-25 on pool nodes Postgres 9.1 xrootd 3.0.4		Upgrade to 2.2-10+ in June 2013 (3rd golden release)
KIT	dCache <ul style="list-style-type: none"> <li>• atlassrm-fzk.gridka.de: 1.9.12-11 (Chimera)</li> <li>• cmssrm-fzk.gridka.de: 1.9.12-17 (Chimera)</li> <li>• lhcbstrm-kit.gridka.de: 1.9.12-24 (Chimera)</li> </ul> xrootd (version 20100510-1509_dbg and 3.2.6)		
NDGF	dCache 2.3 (Chimera) on core servers. Mix of 2.3 and 2.2 versions on pool nodes.		
NL-T1	dCache 2.2.7 (Chimera) (SURFsara), DPM 1.8.6 (NIKHEF)		
PIC	dCache 1.9.12-20 (Chimera) - doors at 1.9.12-23		Head nodes to be migrated to 1.9.12-23 by 10th of June

Site	Status	Recent changes	Planned changes
RAL	CASTOR 2.1.12-10 2.1.13-9 (tape servers) SRM 2.11-1		Upgrading all instances to CASTOR 2.1.13-9 by end of June.
TRIUMF	dCache 2.2.10(Chimera)		

## FTS deployment

Site	Version	Recent changes	Planned changes
CERN	2.2.8 - transfer-fts-3.7.12-1		
ASGC	2.2.8 - transfer-fts-3.7.12-1		
BNL	2.2.8 - transfer-fts-3.7.10-1	Upgrade Oracle to 11g	None
CNAF	2.2.8 - transfer-fts-3.7.12-1		
FNAL	2.2.8 - transfer-fts-3.7.12-1		
IN2P3	2.2.8 - transfer-fts-3.7.12-1		
KIT	2.2.8 - transfer-fts-3.7.12-1		
NDGF	2.2.8 - transfer-fts-3.7.12-1		
NL-T1	2.2.8 - transfer-fts-3.7.12-1		
PIC	2.2.8 - transfer-fts-3.7.12-1		
RAL	2.2.8 - transfer-fts-3.7.12-1		
TRIUMF	2.2.8 - transfer-fts-3.7.12-1		

## LFC deployment

Site	Version	OS, distribution	Backend	WLCG VOs	Upgrade plans
BNL	1.8.0-1 for T1 and 1.8.3.1-1 for US T2s	SL5, gLite	Oracle	ATLAS	Oracle upgraded to 11 on May 28, 2013
CERN	1.8.6-1	SLC6, EMI2	Oracle 11	ATLAS, LHCb, OPS, ATLAS Xroot federations	

## Other site news

- Organized a workshop on May 7th/8th with Tier-0 and Tier-1 sites supporting CMS to plan the separation between disk and tape endpoints
  - https://indico.cern.ch/conferenceDisplay.py?confId=249032
  - Site review
    - CERN - CASTOR for tape, EOS for disk - CMS completing the migration of all Tier-0 workflows to use EOS
    - RAL - currently using CASTOR both for tape and disk with separate namespaces; evaluating different technologies for disk
    - CNAF - solution in StoRM under study: evaluating separate namespaces vs. use of extended attributes
    - KIT, PIC, IN2P3 - new functionality planned in dCache 2.6 to handle disk<-->tape "transfers" with single dCache instance, discussed with dCache developers
    - FNAL - deploy two separate dCache instances for disk and for tape + EOS for user data

About CNAF, Ian comments that an advantage of having separate namespaces for disk and tape data is to make xrootd access easier. Nicolò will discuss with the StoRM developers.

Alessandro is very interested in the details of the disk-tape separation and will discuss it with CMS offline.

## Data management provider news

### Machine/job features (S. Roiser)

For details, see the slides[☞](#).

LHCb would be very interested to use the mechanism devised by the HEPIX virtualisation group and implemented at CERN for LSF and at NIKHEF for Torque/Maui to let jobs access information about the node and the job constraints via local files. Ideally it should be available at all sites, which implies writing implementations for all batch systems in WLCG.

Distinct advantages with respect to information published in the BDII are:

- information specific to the job itself (e.g. remaining wallclock/CPU time before batch system termination, etc.)
- information specific to the node, which e.g. allows for heterogeneous clusters, a more precise HS06 calculation and to let a pilot know how many cores it can use

Alessandra points out that this is very interesting also for the sites.

Ian says that CMS would like a seamless transition to multicore jobs without the need for dedicated multicore queues and this mechanism looks reasonable; the sooner it is available, the better.

Simone confirms the interest of ATLAS and explains the slight loss of momentum in the development needed to use this information by the fact that the first specification was rather volatile.

Stefan reports that according to Ulrich the current specification is stable enough, though changes can be done if needed by the experiments.

It is concluded that the experiments interested in this should carefully check the specification to verify if it is adequate and report for the next meeting. Also, it should be checked if Ulrich is still responsible for this mechanism.

## Experiment operations review and plans

### ALICE (M. Litmaath)

- IN2P3: 110 TB added through 2 additional xrootd servers, thanks!
- CVMFS
  - ◆ CVMFS repository for ALICE has been moved to production infrastructure in IT.
  - ◆ Replication to stratum 1 sites is in progress.
  - ◆ ALICE will proceed with tests on the OpenStack Ibex infrastructure at CERN.
  - ◆ Once it is confirmed that production jobs run fine, ALICE will start configuring some sites in this mode.

### ATLAS (A. Di Girolamo, I. Ueda)

- ATLAS Technical Interchange Meeting held in Tokyo 15-17 May. Main topics: Networks in Production and Analysis workflows, Production system evolution, and Leadership and High Performance Computing. A detailed summary will be given at the ATLAS Software & Computing week (10-14 June).

- Reprocessing of 2012 period-A data has started. First slice tests is being performed. From the infrastructure point of view, respect to the previous campaign, this campaign will use the Frontier/Squid infrastructure to get the conditions data instead of the data replicated on HOTDISK.
- WebDAV access: in order to provide consistent requirements to sites, ATLAS would like to get, from the other experiment interested in the WebDAV access, their experts contact names to organize a technical discussion. (ATLAS plan/request in WLCG Operations planning March 2013 [↗](#))

## CMS (C. Wissing)

- Processing plans/status
  - ◆ Running upgrade MC productions with high prio
  - ◆ Reprocessing of 2011 data/MC still in preparation/validation
- Phedex DBparam Update
  - ◆ Went fine on Tuesday (is this true?)
  - ◆ Big thanks to all CMS site contacts and admins
- WLCG Squid Monitoring Upddate
  - ◆ All T1 sites updated
  - ◆ Two T2 sites still pending (remindes incl. GGUS tickets sent)
  - ◆ Still 18 T3 sites pending
- xrootd federation deployment
  - ◆ 39 Tier-2 support 'fallback'
  - ◆ 32 Tier-2 publish their SE into a xrootd federation
- CVMFS deployment
  - ◆ Detailed CMS statistics in the TF meeting
  - ◆ Summary to WLCG Operation Coordination in next meeting

## LHCb (J. Closier, S. Roiser)

Main activities : MC simulation and stripping. RAL is very efficient and IN2P3 and GRIDKA are suffering of lot of staging requests.

- IN2P3: problem of disk cache for the tape system. and not optimal configuration to gring the data from the tape to the disk (to optimize as discuss 2 years ago with IN2P3 they could get rid of the "export" step).
- CERN : stratum 0 issue but LHCb has not been to much affected.
- GRIDKA : is the tapeset are properly defined for the LHCbtape ?
- SARA: lost one tape. (namespace clean too fast before we were able to recover)

## Task Force reports

### SL6 (A. Forti)

- CERN
  - ◆ Update on lxplus SLC6: we have doubled the number of nodes, 66 now, 8 Gb memory each, 4 cores.
    - ◇ The intermittent login failures to lxplus.cern.ch are happening with less frequency, thanks to some workarounds put in place. The issue is being looked up by the sssd developers.
  - ◆ On lxbatch SLC6: 300 24-core nodes and 304 8-core nodes, in total 8K slots.

- BNL and Nikhef have gone online bringing the total of migrated T1 to 6.
  - ◆ IN2P3-CC is the next one setting in production 25% of resources by the 11th of June
  - ◆ CNAF is waiting for a delivery of new hardware in mid-June to have a clearer plan to upgrade but it looks they'll go for an early summer upgrade.
- Introduced summary stats to the SL6deployment page based on what is reported on the page. This week we have
  - ◆ Total number of Tier1s Done: 6/15 (ALICE 4/9, ATLAS 4/12, CMS 2/9, LHCb 3/8)  
Total number of site in test/tested or Done: 13
  - ◆ Total number of Tier2s Done: 20/131 (Alice 8/41, Atlas 7/91, CMS 11/65, LHCb 4/45)  
Total number of site in test/tested or Done: 47
- Also started to get statistics from the BDII to compare what is published with what is merely in testing. Although this still is not complete because not all sites are in the BDII and the BDII cannot say if the queues are in testing or not it still helps knowing what people are doing in Europe and some other place.
  - ◆ Here are the (7) T1s publishing IN2P3-CC, JINR-T1, NIKHEF-ELPROD, pic, RAL-LCG2, SARA-MATRIX, TRIUMF-LCG2
  - ◆ While there are 29 T2s.
- Updated the procedures section with extra link for ATLAS to upgrade work arounds and problems
  - ◆ So far reported kernel vulnerability which was in the EGI Alert 2013-05-14 [\[7\]](#) that sites have received. This affects of course any site and is followed up by the EGI CSIRTs but it is a useful reminder for sites either to apply the patches or to upgrade to the latest kernel.
  - ◆ On SL6 there is a problem with flags reported by the open() call when files exist and this is causing some ATLAS applications to fail when they access CVMFS files . The workaround is to mount CVMFS in rw mode. This will become default in CVMFS version 2.1.11 but needs to be changed by sites if they use 2.0.x or <2.1.11. This has been reported in INC299258 [\[8\]](#). It is in test in some sites and BNL has the work around in production, they haven't reported any side effects.
  - ◆ An additional problem already reported about excessive number of file descriptors left on the system is still under investigation because not all sites are reporting it.

## gLExec (M. Litmaath)

- May 14 Management Board minutes [\[9\]](#)
- gLExec deployment tracking page
  - ◆ formatting still expected to change
- OSG will look into their sites

Maria G. asks if there is a tentative deployment schedule. Maarten says that it has not yet been discussed and he plans to have more information for the next meeting. Oliver G. adds that CMS has a SAM test for gLExec but so far nobody was tasked with following up on gLExec issues.

## SHA-2 migration (M. Litmaath)

- SHA-2 testing instructions. Anybody with a CERN account is entitled to get a SHA-2 certificate from CERN.
- the SHA-2 CERN CA will be included in the June IGTF release
- the SHA-2 VOMS server now is accessible also from outside CERN

Maarten adds that there is no recent news from IT-PES on the timescale for migrating from VOMRS (not SHA-2 compliant) to VOMS-admin.

## FTS-3 (M. Salichos, A. Di Girolamo)

- FTS demo held 22nd May.
  - ◆ deployment model: ATLAS reported that it would be important for the experiment to have FTS3 instances being able to share configuration and scheduling information between themselves. Various Tier1s reported too their interest in having FTS3 instances with this functionality, so that they can run their own instance of FTS3. This will be investigated by the developers.
  - ◆ production services: under discussion with CERN and RAL the possibility to have production quality FTS3 services.

Maria G. asks when the TF will expose a deployment schedule. Alessandro answers that they are waiting for the production nodes where to deploy the production service and for the first official FTS-3 release. Oliver adds that for the former, there is no timescale from IT-PES and for the latter, it should take approximately one month.

Maria urges the TF to create a dedicated twiki as it makes it easier for people not in the TF to track the progress.

## perfSONAR (S. Campana)

Our testing of RC4 has looked good in the US and CERN. Andy Lake from ESnet reported that RC4 looks good so far. If no problematic report comes on RC4 by the end of the week the release process will start. We are close to having v3.3 for production.

## WLCG monitoring support unit (M. Babik)

For the details, see the slides [☞](#).

IT-SDC proposes to unify SAM and Dashboard support in GGUS to expose to users a single entry point to WLCG monitoring. In addition, they will establish common WLCG monitoring shifts for support.

Peter asks if this will decouple support for WLCG and generic SAM support; Marian explains that indeed the current SAM support unit will be eventually phased out, but EGI uses a 2nd line of support and in agreement with EGI tickets will be moved from it to the new SU.

The timeline for the change is to complete it by this autumn. Maria D. warns that due to the GGUS release schedule, the creation of the new SU will be possible only in July or from September. Marian proposes to use the July release and deal with the interfacing with SNOW at a later time.

Alessandra asks if this change will also replace Savannah with GGUS or JIRA for Dashboard issues. Pablo and Maarten explain that GGUS will indeed be the official support entry point, and whether the WLCG monitoring team will use JIRA internally is irrelevant to the users.

It is agreed that by the next meeting (June 20) people can give feedback on the proposal.

## AOB

- Important GGUS release next Wednesday 2013/06/05, especially due to the amount of Support Units being decommissioned. Check the dev. items' list <http://bit.ly/14Jhw0C> [☞](#) for details

In order to realign the meeting dates with the 1st and 3rd Thursday of each month, these are the agreed dates for the June meetings:



- June 6: virtual meeting (reports to be added to the minutes twiki)
- June 20: real meeting, as usual

## Action list

1. Add perfSONAR to the baseline versions list as soon as the latest Release Candidate becomes production-ready.
2. Build a list by experiment of the Tier-2's that need to upgrade dCache to 2.2
  - ◆ **done** for ATLAS: list updated on 17-05-2013
  - ◆ **done** for CMS: list updated on 03-05-2013
  - ◆ not applicable to LHCb, nor to ALICE
  - ◆ Maarten: EGI and OSG will track sites that need to upgrade away from unsupported releases.
3. Inform sites that they need to install the latest Frontier/squid RPM by May at the latest ( **done** for CMS and ATLAS, status monitored🔗)
4. Maarten will look into SHA-2 testing by the experiments when experts can obtain VOMS proxies for their VO.
  - ◆ started, see SHA-2 readiness testing
5. Tracking tools TF members who own savannah projects to list them and submit them to the savannah and jira developers if they wish to migrate them to jira. AndreaV and MariaD to report on their experience from the migration of their own savannah trackers.
6. CMS to appoint a contact point for Savannah:131565🔗 and dependent dev. items for the savannah-ggus bridge replacement by a GGUS-only solution tailored to CMS ops needs.
7. Investigate how to separate Disk and Tape services in GOCDB
8. Update the CVMFS baseline version to the latest version **DONE**
9. Agree with IT-PES on a clear timeline to migrate OPS and the LHC VOs from VOMRS to VOMS-Admin
10. For the experiments to give feedback on the machine/job information specifications
11. Experiments interested in using WebDAV should contact ATLAS to organise a common discussion
12. Create a twiki for the FTS-3 task force
13. Give feedback to the IT-SDC proposal to unify WLCG monitoring support in GGUS

## Chat room comments

-- AndreaSciaba - 28-May-2013

---

This topic: LCG > WLCGOpsMinutes130530  
 Topic revision: r30 - 2013-06-11 - AndreaSciaba



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