

Table of Contents

WLCG Operations Coordination Minutes - 7th March 2013.....	1
Agenda.....	1
Attendance.....	1
News (M. Girone).....	1
Middleware news and baseline versions (N. Magini).....	1
dCache.....	1
Tier-1 Grid services.....	2
Storage deployment.....	2
FTS deployment.....	3
LFC deployment.....	3
Other site news.....	3
Data management provider news.....	4
UMD release plans (P. Solagna).....	4
Experiment operations review and plans.....	4
ALICE (M. Litmaath).....	4
ATLAS (I. Ueda).....	4
CMS (I. Fisk).....	5
LHCb (J. Closier).....	5
Enforcement of policy on personal data retention (P. Solagna).....	5
Task Force reports.....	6
CVMFS.....	6
gLExec (M. Litmaath).....	6
perfSONAR (S. McKee).....	6
SHA-2 migration (M. Litmaath).....	6
Middleware deployment.....	7
FTS 3 integration and deployment (N. Magini).....	7
Xrootd deployment.....	7
Tracking tools (M. Dimou).....	7
SL6 migration task force (A. Forti).....	7
News from other WLCG working groups.....	8
AOB.....	8
Action list.....	8
Chat room comments.....	9

WLCG Operations Coordination Minutes - 7th March 2013

Agenda

- <http://indico.cern.ch/conferenceDisplay.py?confId=239108>

Attendance

- Local: Maria Girone (chair), Andrea Sciabà (secretary), Maria Dimou, Ian Fisk, Maarten Litmaath, Massimo Lamanna, Xavier Espinal, Felix Lee, Andrea Valassi, Simone Campana, Michail Salichos, Nicolò Magini, Ikuo Ueda, Maite Barroso Lopez, Luca Mascetti, Alessandro Di Girolamo.
- Remote: Alessandra Forti, Renaud Vernet, Massimo Sgaravatto, Joel Closier, Shawn McKee, Christoph Wissing, Peter Solagna, Stephen Burke, Daniele Bonacorsi, Di Qing, Gareth Smith, Jeremy Coles, Burt Holzman, Ron Trompert.
- Apologies: Ian Collier

News (M. Girone)

From today we have a standing agenda item for the first meeting of each month about news from EGI, in particular about UMD updates.

From this week the daily WLCG operations meeting has become twice a week: Maria D. (SCOD this week) reports that the meeting duration did not increase.

Middleware news and baseline versions (N. Magini)

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

Minor changes: there is a patch for the WMS client that fixes the known issue in the EMI-2 UI causing a fraction of WMS jobs to abort. Sites with an EMI-2 UI are encouraged to upgrade.

dCache

As dCache 1.9.* is reaching end of support at the end of April, all Tier-1 and Tier-2 sites should consider upgrading to the new golden release (2.2.*) as soon as possible; EGI (according to a policy agreed with WLCG) has a hard deadline for decommissioning dCache (as any other service) one month after end of support, which for dCache would be on May 31. For Tier-1 sites this is not trivial because they need to test their interfaces to the tape backends; only NL-T1 and NDGF have already moved to 2.2.

Maria G. asks what concrete action might be proposed; Nicolò suggests to open GGUS tickets to the Tier-1's to know about issues, schedules and plans; concerning Tier-2's, they should have less problems in upgrading and a more aggressive schedule should be possible.

Maarten thinks that for Tier-2's this is the usual business of upgrading services becoming obsolete and this is normally taken care of by EGI, so WLCG operations coordination should mainly worry about Tier-1 sites. Still, in the past in similar situations sites upgraded dCache by themselves without need for a strong orchestration.

Maria G. suggests to use the Tier-1 storage service table in the minutes for sites to communicate their upgrade

plans. Concerning Tier-2's, it is easy to build a list of Tier-2's needing to upgrade for each VO using the BDII. It is decided that such lists will be generated and sites informed.

Joel says that it is not clear which version of the lcg_util and the DPM clients should be installed, as it is not explicitly said in the baseline versions table. It is decided that this information will be added to the table.

It is stressed again that the baseline versions table lists the versions that for WLCG must **at least** be installed, but the general rule is that the latest versions are even better, and any exception to the rule will be made explicit.

Tier-1 Grid services

Storage deployment

Site	Status	Recent changes	Planned changes
CERN	CASTOR: 2.1.13-9 being deployed next week for all experiments (already in production on CASTORPUBLIC) / SRM-2.11 for all instances. EOS: ALICE (EOS 0.2.20 / xrootd 3.2.5) 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)		CASTOR: Close the possibility to update files (files will be immutable) - feature barely used (0.01/million) CASTOR: root protocol will be phased out, barely used (will contact users still using it) EOS: upgrades to 0.2.29 for ALICE will be scheduled in agreement with the experiment New Bestman2 release to be tested, packaged and deployed (bugfix for sha2)
ASGC	CASTOR 2.1.13-9 CASTOR SRM 2.11-2 DPM 1.8.6-1 xrootd 3.2.7-1	Feb 25th: unscheduled downtime to OPN links due to fire Feb 26th: scheduled downtime for CASTOR 2.1.13 and DPM 1.8.6 upgrades	None
BNL	dCache 1.9.12.10 (Chimera, Postgres 9 w/ hot backup) http (aria2c) and xrootd/Scalla on each pool	None	None
CNAF	StoRM 1.8.1 (Atlas, CMS, LHCb)		
FNAL	dCache 1.9.5-23 (PNFS, postgres 8 with backup, distributed SRM) httpd=2.2.3 Scalla xrootd 2.9.7/3.2.4-1.osg Oracle Lustre 1.8.6 EOS 0.2.22-4/xrootd 3.2.4-1.osg with Bestman 2.2.2.0.10		
IN2P3	dCache 1.9.12-16 (Chimera) on SL6 core servers and 1.9.12-24 on pool nodes Postgres 9.1 xrootd 3.0.4		
KIT	dCache		

Site	Status	Recent changes	Planned changes
	<ul style="list-style-type: none"> • atlasrm-fzk.gridka.de: 1.9.12-11 (Chimera) • cmssrm-fzk.gridka.de: 1.9.12-17 (Chimera) • lhcbarm-kit.gridka.de: 1.9.12-24 (Chimera) 		
	xrootd (version 20100510-1509_dbg)		
NDGF	dCache 2.3 (Chimera) on core servers. Mix of 2.3 and 2.2 versions on pool nodes.		
NL-T1	dCache 2.2.4 (Chimera) (SARA), DPM 1.8.2 (NIKHEF)		Mon 11th network maintenance
PIC	dCache 1.9.12-20 (Chimera) - doors at 1.9.12-23		
RAL	CASTOR 2.1.12-10 2.1.13-9 (tape servers) SRM 2.11-1	Upgraded tape servers to 2.1.13-9	Upgrading WN CASTOR clients, NS and stagers to 2.1.13-9
TRIUMF	dCache 1.9.12-19(Chimera)		doing further tests with java7, may upgrade to dcache 2.2.8 in April

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	None	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	None
CERN	1.8.6-1	SLC6, EMI2	Oracle 11	ATLAS, LHCb, OPS, ATLAS Xroot federations	

Other site news

Xavi adds that CERN plans to drop the possibility to update files (as is almost never used) and the support for the root (**not** xroot!) protocol in CASTOR.

Data management provider news

UMD release plans (P. Solagna)

In February there were updates in UMD 1 (mostly security fixes) and in UMD 2, clearing up the queue of updates from EMI-2 apart from some released in January, which should be released only after the first EMI-3 release makes it to UMD (end of April, sooner if possible). EGI proposed a prioritisation of EMI-3 products (high priority and medium priority) for the first UMD 3.0 release: it is expected that all high priority and most medium priority will make it; feedback can be provided until March 15. EGI proposes to create UMD 3.0 for the EMI-3 products, which helps in keeping them separated from EMI-2 and makes its eventual decommissioning easier. The only negative effect would be the need for sites to change the repository configuration for new services. Yet another UMD major release will be likely needed after the end of EMI/IGE. Again, feedback is welcome before March 15.

It is agreed that the creation of UMD 3.0 is the way to go.

Experiment operations review and plans

ALICE (M. Litmaath)

- Central services: on Feb 25 the AliEn catalogue DB was moved to a new, more powerful machine to sustain its steady growth - 10^9 entries were reached on Feb 21.
- KISTI
 - ◆ GLORIAD-CERN network maintenance on Feb 27 made the site unusable for 1 day instead of the expected 1h downtime
 - ◆ main disk SE was unstable for 3 days, OK again since early March 4
- KIT
 - ◆ disk SE was unstable for 3 days, leading to high load on the central firewall when jobs failed over to remote SEs; fixed since March 4 late afternoon

ATLAS (I. Ueda)

Activities:

- The important winter conference started == the majority of the very important production/analysis jobs are done.
 - ◆ starting a series of small scale (re)processing of real data (will trigger staging of data from T1 tapes)
 - ◆ starting some important production for later conferences
- ATLAS has been successfully commissioning the RU-T1 prototype (RRC-KI-T1) in ATLAS systems.
 - ◆ RRC-KI-T1 has been included in ATLAS DDM.
 - ◇ Transfer throughput, after first days of commissioning, is now up to 250MB/s sustained over a day with efficiency above 90%
 - ◇ FTS3 is used to send data to RRC-KI-T1. CERN FTS228 used to send data back from RRC-KI-T1 to CERN.
 - ◆ RRC-KI-T1 has been included in Panda. HammerCloud has been used to test (and stress test) the functionality of the production queue.
 - ◇ up to 1250 parallel jobs
 - ◇ tasks to validate the ATLAS delay stream have been submitted: some of them are now successfully finishing.
- we believe that the ATLAS experiences to commission the RRC-KI-T1 can be useful also to the other Experiments

Issues:

- SAM tests results delayed: ATLAS observed a problem in getting the latest results of SAM tests (input for Storage-Area-Automatic-Blacklisting system - SAAB) Monday and Tuesday 4 and 5 March.
SAM Team feedback:
 - ◆ Monday: SAM update-20 intervention in the morning which caused delay up to Tuesday morning
 - ◆ Tuesday: performances issue with one of status computation procedures (related to LCGR intervention) - recovered overnight

Concerning the SAM problem, Maria G. encourages to use a ticket next time for better tracking.

CMS (I. Fisk)

- CMS is progressing well re-reconstructing the 2012 data
- Because all the T1 resources are taken for data re-reconstruction, CMS tested and is moving MC digitization and reconstruction workflows to big T2 sites, reading the input simulated events via xrootd
 - ◆ We started with the US T2 sites and will expand to German and Italian T2 sites soon
- Plan to reconfigure was generally approved and being worked on
 - ◆ Some delay in reconfiguring LSF which we could use right now better for re-processing
 - ◆ Castor disk pools are being cleaned and we will ask soon to move them to EOS
- HLT cloud commissioning is progressing after the network reconfiguration
- Site issues: IN2P3
 - ◆ After having problems with direct dcap reads and switching to xrootd reads, also hitting limitations in xrootd file access (bug was filed to dCache team)
 - ◆ Proposed several solutions involving fallback to local xrootd or even reading RAW data from CERN via xrootd to help with the situation
- ASGC:
 - ◆ evacuating custodial data and MC samples from ASGC:
 - ◇ MC will go to CERN: transfers setup, checks with CERN IT for tape space complete, but transfers not approved yet
 - ◇ Data will be distributed amongst the other T1 sites to retain two copies on tape

It is clarified that also T2_TW_Taiwan will be shut down and that, according to the MoU agreements, CMS has 18 months to copy the custodial data out of ASGC, even if there will not be any CPU resources available much sooner than then. Nicolò and Ian add that realistically CMS will not need more than 1-2 months.

LHCb (J. Closier)

- ask Grid Middleware to provide the libtool-ltdt dependency on LCG AA area. (GGUS:91882[a](#)).

Enforcement of policy on personal data retention (P. Solagna)

EGI proposes to extend the data retention policy for individual accounting records containing personal user information from 12 to 18 months from July 1st. The WLCG management and the WLCG security experts are well aware of this and did not make any objection before the deadline (tomorrow).

Task Force reports

CVMFS

No report.

gLExec (M. Litmaath)

- CMS
 - ◆ gLExec being used at 20 EGI sites, 15 OSG sites
 - ◆ status [↗](#) of gLExec tests (193 CE on March 6)
 - ◇ only the CEs tested successfully: here [↗](#) (93 on March 6)
- LHCb The verification of the functionality in DIRAC should start next week.

Maria G. asks about the status of a plan with milestones. Maarten answers that he needs to know what timelines can work for the experiments. For CMS, Ian proposes July 1st to have at least 90% of sites with gLExec enabled and he will bring this up at the next CMS computing operations meeting.

Christoph asks if the *log-only* mode is considered acceptable. It is agreed that the *setuid* mode is the preferred solution but for those cases where it is not possible (there should be only a few) the *log-only* mode is still better than nothing.

perfSONAR (S. McKee)

- The basic services provided by perfSONAR-PS have Nagios plugins available to verify proper operation. We are preparing SAM/Nagios tests to test the hosts (this was introduced by Alessandra last time)
- The mesh functionality includes DISJOINT meshes where GroupA member can test to all GroupB members but there is no testing within GroupA or GroupB. This is being tested now
 - ◆ Additionally we have the ability to include JSON files at the mesh definition level. Working on best practice uses.
- We are testing the new release of perfSONAR 3.3RC1 and, soon, RC2. Expect final release of 3.3 this month, followed by a big push to get this in place within WCLG.

SHA-2 migration (M. Litmaath)

- Still waiting for the new CERN CA, hopefully next week

The biggest concern is for those experiment services that use components known to have problems with SHA-2, like Gridsite and the Trust Manager. Still, these services are not many.

Generic links:

- Jan pre-GDB update [↗](#)
 - ◆ Jan pre-GDB minutes
- RFC proxy and SHA-2 signature support in WLCG middleware
- EGI links:
 - ◆ <https://documents.egi.eu/public/ShowDocument?docid=1291> [↗](#)
 - ◆ https://wiki.egi.eu/wiki/EGI-JRA1_SHA2_Readiness [↗](#)
- OSG links:
 - ◆ <https://twiki.grid.iu.edu/bin/view/SoftwareTeam/Sha2Support> [↗](#)
 - ◆ <https://twiki.grid.iu.edu/bin/view/Security/HashAlgorithms> [↗](#)

Middleware deployment

FTS 3 integration and deployment (N. Magini)

- Stress testing and scalability testing results will be presented in the next WLCG FTS3 task force meeting and demo
- Integration testing: ATLAS used FTS3 pilot to populate new Russian T1 with production data; CMS T2 using FTS3 to import test data from all sites to verify effect of auto-tuning
- Bulk SRM BringOnline operation implementation has been completed, will be installed in the pilot sometime next week
- Retry failed transfers logic has been completed, categorization of the recoverable/non recoverable errors will be discussed in the next WLCG FTS3 task force meeting
- RESTful interface for transfer submission and status retrieval demonstrated

Maria G. asks what is the status of the deployment schedule. Nicolò and Alessandro answer that in the next TF meeting the results of the scale tests will be discussed and given that they are needed to make a deployment plan, they hope to be able to have a proposal by April.

Alessandro adds that there is a show-stopper for StoRM sites that have not upgraded to the EMI-2 1.10 version, which is mandatory for FTS-3 to work.

Xrootd deployment

No report.

Tracking tools (M. Dimou)

- According to the last minutes https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsMinutes130221#Tracking_tools we can only report on savannah-jira migration progress after Easter when we shall define with the savannah developers the field mapping for the GGUS dev. tracker (migration example for the community).
- We need a CMS contact to continue the savannah-ggus bridge migration to a GGUS only solution. Implementation solutions were discussed on 2012/12/04 and are documented in Savannah:131565 [↗](#) and Savannah:134411 [↗](#), Savannah:134413 [↗](#), Savannah:134415 [↗](#), Savannah:134416 [↗](#).

SL6 migration task force (A. Forti)

People who have joined:

- T0: Helge Meinhard, Steve Traylen
- T1: Ian Collier (RAL), Di Qing (TRIUMF), Burt Holzman (FNAL)
- T2: Alessandra Forti (Manchester), Shawn McKee (AGLT2), Raul Lopez (Brunel), Alessandra Doria (Napoli)
- Atlas: Simone Campana (ops), Alessandro De Salvo (ops), Rod Walker (ops), Ikuo Ueda (ops), Emil Obreshkov (sw librarian)
- CMS: Christoph Wissing (desy), Giulio Eulisse (sw librarian), Oliver Gutsch (computing operations), Brian Bockelman (grid and xrootd expert)
- Lhcb: Stefan Roiser (ops), Ben Couturier (sw librarian), Joel Closier (grid expert)
- Alice: Maarten Litmath (ops), Latchezar Betev (offline coord)
- EGI: Tiziana Ferrari, Peter Solagna
- SL6 tarball: Matt Doidge (Lancaster)
- IT/ES: Andrea Valassi

Mailing list: egroup_wlcg-ops-coord-tf-sl6-migration.

Twiki page: <https://twiki.cern.ch/twiki/bin/view/LCG/SL6Migration>.

First points I collected in the discussions already had either in this forum or at the GDB

1. Understand each experiment sites status, i.e. Alice has already some sites on SL6, Atlas has 1, CMS? Lhcb?
2. Put together the documentation necessary for sites: some twiki pages already exist are there others? Are they all visible to the external world or do they require special access?
3. Test HEPOS_libs on external sites not using SLC6
4. Do we need test queues at sites? How can the upgrade be done? Atlas needs the SL6 nodes separated from the SL5 ones, I'm told CMS doesn't care because the pilot can test (is that true?), Lhcb? Alice?
5. What communication channels should we use to help sites move? Experiments? Tickets?
6. Do we need to follow every site?
7. Do we need to coordinate T1s? Atlas doesn't want all T1s going at the same time for example what are the other experiments thinking?
8. OSG sites? At the moment the representation is mostly EU centric + Canada.
 - ◆ FNAL and AGLT2 are represented now.
9. Do we need a target date? If yes it is important it either is far away from EMI-3 migration or it coincides with it. EMI-3 migration is by 30 April 2014.
10. ~~Issue with Lxplus migration timeline raised by CMS~~. This has been solved.

Concerning the proposed target date for the migration to complete, it is agreed that September is too aggressive due to the summer holidays and October is more realistic. The expectation is that most sites will have migrated by then, and the rest of autumn will be needed for the "tails". Alessandra adds that many sites are eager to start now; more than half of the US-CMS sites have already migrated and US-ATLAS is pushing as well.

News from other WLCG working groups

AOB

Maria G. reminds that the next will be a planning meeting, in two weeks.

Action list

- Build a list by experiment of the Tier-2's that need to upgrade dCache to 2.2.
- Inform sites that they need to install the latest Frontier/squid RPM by April at the latest.
- Inform CMS sites that they must configure a queue with a length of at least 48 hours, if they have not done it already. **DONE**
- Inform CMS DPM sites that they should enable the xrootd interface. **DONE**
- Maarten will look into SHA-2 testing by the experiments when the new CERN CA has become available.
- MariaD will convey savannah developers OliverK's idea to place a banner on every savannah ticket warning about the switch off date. **DONE**: Savannah:134651#comment14
- 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.

Chat room comments

Meeting chat room comments

-- AndreaSciaba - 05-Mar-2013

This topic: LCG > WLCGOpsMinutes130307

Topic revision: r28 - 2013-03-11 - NicoloMagini



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](#)