

# Table of Contents

<b>WLCG Operations Coordination Minutes - 19 September 2013</b> .....	<b>1</b>
Agenda.....	1
Attendance.....	1
News.....	1
Middleware news and baseline versions.....	1
Tier-1 Grid services.....	2
Storage deployment.....	2
FTS deployment.....	3
LFC deployment.....	4
Other site news.....	4
Data management provider news.....	4
Experiment operations review and plans.....	4
ALICE.....	4
ATLAS.....	5
CMS.....	5
LHCb.....	6
Task Force reports.....	6
SHA-2 migration.....	6
Machine/Job Features.....	7
CVMFS.....	7
perfSONAR.....	7
FTS-3.....	7
Tracking tools.....	8
IPv6.....	8
XrootD Deployment.....	8
gLExec.....	9
AOB.....	9
Action list.....	9

# WLCG Operations Coordination Minutes - 19 September 2013

## Agenda

<https://indico.cern.ch/conferenceDisplay.py?confId=263202>

## Attendance

- Local: Andrea Sciabà (secretary), Ken Bloom, Stefan Roiser, Oliver Keeble, Simone Campana, Nicolò Magini, Andrea Valassi, Felix Lee, Marten Litmath, Mite Barroso Lopez, Alessandro Di Girolamo, Jan Iven
- Remote: Alessandra Forti (chair), Alessandro Cavalli, Frederique Chollet, Helge Meinhard, Christoph Wissing, Pavel Weber, Massimo Sgaravatto, Daniele Bonacorsi, Isidro Gonzalez Caballero, Alessandra Doria, Gareth Smith, Di Qing, Michael Ernst

## News

Marten and Mite introduce the proposal to start the WMS decommissioning task force, already foreseen since long, because of CERN's plans to decommission the WMS service as soon as practical. The residual usage comes from SAM, some small fraction of the LHC experiment jobs and smaller VOs (ILC, GEANT4). The task force will start by analysing the logs to determine exactly who are the users, and then to discuss alternatives with them. A twiki will soon be created, and other volunteers are welcome.

About the ops tests, Marten explains that WLCG will stop looking at them from January, but they will still be run by EGI for their monitoring.

The task force is approved and its coordinators will be Marten and Mite.

## Middleware news and baseline versions

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

Highlights:

- minor change for BDI I
- new baseline version for dCache is 2.2.17, introducing SHA-2 support
- new baseline version for Stormis 1.11.2, in EM-3. It fixes some known issues (also related to FTS 3) and supports SHA-2. Sites should upgrade

Marten adds that earlier this week a critical bug was discovered in the top BDI I. Sites have been notified by EGI. A patch release is expected next week (not this week, because an easy workaround is available); still, upgrading to the current latest version is advisable.

## Ti er - 1 G r i d s e r v i c e s

### St o r a g e d e p l o y m e n t

Si t e	St a t u s	Re c e n t c h a n g e s	Pl a n n e d c h a n g e s
CERN	CASTOR: v2.1.13-9-2 and SRM 2.11-2 for all instances EOS: ALICE (EOS 0.2.37 / xrootd 3.2.8) ATLAS (EOS 0.2.38 / xrootd 3.2.8 / BeStMan2-2.2.2) CMS (EOS 0.2.38 / xrootd 3.2.8 / BeStMan2-2.2.2) LHCb (EOS 0.2.29 / xrootd 3.2.7 / BeStMan2-2.2.2)	2013-09-18: EOSCMS updated to 0.2.38 (quota issue)	
ASGC	CASTOR 2.1.13-9 CASTOR SRM 2.11-2 DPM 1.8.6-1 xrootd 3.2.7-1	None	None
BNL	dCache 2.2.10 (Chimera, Postgres 9 w/ hot backup) http (aria2c) and xrootd/Scalla on each pool	None	None
CNAF	StoRM 1.11.2 emi3 (Atlas, LHCb) StoRM 1.8.1 (CMS)		In 2 weeks upgrade to emi3 for Alice/CMS
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.slc Oracle Lustre 1.8.6 EOS 0.3.1-5/xrootd 3.3.3-1.slc5 with Bestman 2.2.2.0.10	EOS in production for volatile data (temporary CMS unmerged files)	
IN2P3	dCache 2.2.12-1 (Chimera) on SL6 core servers and 2.2.13-1 on pool nodes Postgres 9.1 xrootd 3.0.4		
KISTI	xrootd v3.2.6 on SL5 for disk pools xrootd 20100510-1509_dbg on SL6 for tape pool dpm 1.8.6		xrootd upgrade foreseen for tape (20100510-1509_dbg -> v3.1.1) in September
KIT	dCache		

	<ul style="list-style-type: none"> <li>atlassrm.fzk.gridka.de: 2.6.5-1</li> <li>cmssrm.fzk.gridka.de: 2.6.5-1</li> <li>lhcbstrm.kit.gridka.de: 2.6.5-1</li> </ul>		alice-tape-se.gridka.de will be upgraded to xrootd 3.2.6 in Sep/Oct. The ATLAS FAX xrootd proxy will be upgraded to 3.3.3 during September.
	<ul style="list-style-type: none"> <li>alice-tape-se.gridka.de 20100510-1509_dbg</li> <li>alice-disk-se.gridka.de 3.2.6</li> <li>ATLAS FAX xrootd proxy 3.3.1-1</li> </ul>		
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 head nodes (Chimera) and doors at 2.2.7 xrootd 3.3.1-1	16/09/2013: migration from 1.9.12-23 to 2.2.7	
RAL	CASTOR 2.1.13-9 2.1.13-9 (tape servers) SRM 2.11-1		
TRIUMF	dCache 2.2.13(chimera), pool/door 2.2.10		

## 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	FTS3 test instance deployed	
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 VO's	Upgrade plans
BNL	1.8.3.1-1 for T1 and US T2s	SL6, gLite	ORACLE 11gR2	ATLAS	None
CERN	1.8.6-1	SLC6, EM2	Oracle 11	ATLAS, LHCb, OPS, ATLAS Xroot federations	

## Other site news

### Data management provider news

DPM 1.8.7 has been released to EPEL.

[https://svnweb.cern.ch/trac/lcgdm/attachment/wiki/Dpm/dpmrelnotes\\_082013-1.txt](https://svnweb.cern.ch/trac/lcgdm/attachment/wiki/Dpm/dpmrelnotes_082013-1.txt)

This is primarily a bugfix release. For WebDAV, it introduces both bugfixes and support for configuration of arbitrary ports, required by Atlas.

A note on the release strategy - the rpms corresponding to this release have all been removed from the EM repositories and the release has been made directly to EPEL. The release is compatible with EM2 and EM3. Sites running EM2 or EM3 DPMs should simply upgrade as normal and the necessary rpms will be taken from EPEL.

## Experiment operations review and plans

### ALICE

- CVMFS
  - ◆ 38 tickets closed, 14 open, 5 done since last meeting
  - ◆ A few more sites have been switched, in particular KIT!
- KIT
  - ◆ Job submission not working Sat morning Aug 31 through Mon morning Sep 2 due to batch system crash
  - ◆ The long-standing instability of the local SE has been **RESOLVED** since 10 days, after switching off a configuration option that looked unrelated!
  - ◆ Successfully switched to **CVMFS** yesterday morning!
- Russian sites

- ◆ On evening Sep 16 the Russian GEANT link was again limited to 100 Mbps
- ◆ Most T2 sites are still closed for ALICE jobs
  - ◇ JINR, PNPI, Troitsk, ITEP, MEPHI
- ◆ 3 sites are operational
  - ◇ RRC-KI, SPbSU, IHEP (much reduced)

## ATLAS

- arcproxy: due to the problems introduced by the new java version distributed with EM-3 needed by e.g. voms-proxy-info, ATLAS has decided to replace its pilot calls with arcproxy ones. The latest problem ([https://ggus.eu/ws/ticket\\_info.php?ticket=97230](https://ggus.eu/ws/ticket_info.php?ticket=97230)) is not easily solvable since java doesn't like the virtual memory being limited. Limiting the memory to keep out some memory hungry jobs is current practice at many sites. The minimum value sites can set the limit to is not predictable as it looks to be connected to the amount of memory+swap. arcproxy is written in C++ and is considered more suitable for the WNs activity.
- webDAV: for DPM as reported in the baseline versions which have been updated, the required version for ATLAS is DPM 1.8.7 .

Marten clarifies that "the new Java version" refers to the new VOMS clients version written in Java instead of C++ (note that EM does not distribute Java).

## CMS

- Operations and Production activities
  - ◆ legacy 7TeV data rereco started on Tuesday
  - ◆ running 8TeV MC GEN-SIM and started 7 TeV legacy MC GEN-SIM
  - ◆ MC Workflow loaded large Gridpack file via Squid infrastructure (700 MB)
    - ◇ Launchpads started to die
    - ◇ Post Mortem ongoing
- Requests/News for sites
  - ◆ Starting with CMSW\_6\_1\_0, the Xrootd file-close monitoring has been implemented as a CMSW framework service. This allows the CERN popularity service to monitor all file accesses done by CMSW applications. It is deactivated by default.
  - ◆ All sites, please activate the monitoring in your site-local.config.xml following this link
- CVMFS Migration
  - ◆ Two more sites moved to CVMFS
  - ◆ 1 Tier-1 and 8 Tier-2 sites missing
- General Items
  - ◆ IPV6 for VMs at CERN
    - ◇ VMs might get only IPv6 address (reference )
    - ◇ Status of the WLCG Operations Coordinations TF?
  - ◆ SAM tests: condor\_g mode, progress?

Marten replies that work on a new SAM job submission probe based on Condor-G will start in November.

## LHCb

- WLCG services
  - ◆ FTS3 put back into production for transfers going in and out of CERN. The submission time is still reported in local time instead of UTC and currently hotfixed within LHCb sw.
- Sites
  - ◆ GRIDKA staging stress test carried out in the last 2 weeks. Good exercise to spot weaknesses in the site and experiment systems (e.g. increase of tape libraries, better monitoring).
    - ◇ If other sites are not confident about their staging performance they are invited to carry out such a test with LHCb
  - ◆ After the router upgrade campaign in the CERN CC many Dirac services needed to be restarted because they lost connections.
- Experiment Activities
  - ◆ Incremental stripping campaign close to start, last tests on the output rates are carried out by end of this week. Likely to (slowly) start the campaign by next week. Tentative duration ~ 8 weeks. Lots of stress on the sites tape systems.
  - ◆ Requested throughput rates for an incremental stripping campaign lasting 60 days

Total	Rate (MB/s)
CERN	50
CNAF	153
GRIDKA	124
IN2P3	134
PI C	39
RAL	111
SARA	104

## Task Force reports

### SHA-2 migration

- EUGridPMA/IGTF okayed the requested delay until Dec 1 before CAs might start making SHA-2 the default (while still supporting SHA-1)
  - ◆ see the minutes [of the Sep 2013 EUGridPMA meeting](#)
- EGI plan a **mandatory decommissioning** of non-SHA-2 compliant service instances by Nov 30
  - ◆ details to be finalized during the Sep 27 Operations Management Board meeting [?](#)
- CERN VOMRS migration
  - ◆ test instance to be made available in the autumn, date not yet agreed
  - ◆ SHA-2 certs can be registered as secondary certs for the time being, as described [here](#)
- ALICE
  - ◆ MnALISA OK

- ◆ soon: switch the VOBOX services of a small site to the use of a SHA-2 proxy
  - ◇ no problems foreseen

Alessandra F. asks if experiments tested the full job submission chain since last meeting. Marten answers that, once the VO central services and the site services are known to work with SHA-2, there shouldn't be too much to worry about, although running the whole chain would be good as final check. ALICE will do it, but they are concerned mostly about the interaction between the VOBOX and some ALICE central services.

Michael adds that in USATLAS they will definitely do vertical tests, and even if BNL is not yet at the latest version of dCache, there will be no problem meeting the December 1st deadline. About CMS, Andrea S. reminds that the full chain has already been tested with SHA-2 user proxies and glexec; using a SHA-2 proxy for the pilot is still to be done.

## Machine/ Job Features

- TF started effectively in September. Two meetings so far used for requirements gathering of the experiments.
  - ◆ Mainly interested in information on number of cores, time left (wall/cpu), machine power (HepSpec06)
  - ◆ All VOs are interested to retrieve the information and in a first instance compare them to their own observations
- For "physical batch systems" we have already implementations, need to work on virtual infrastructures
- Decision to propose a data structure (e.g. json) to communicate the information, which will be the same for physical batch systems and virtual infrastructures

## CVMFS

- Ongoing good progress in the ALICE deployment campaign, 14 sites left (5 done). Out of those 8 stated they will deploy in Sept.

## perf SONAR

See the slides [☞](#).

All sites are invited to install (or upgrade to) version 3.3.1. A broadcast will be sent.

## FTS-3

- IT-PES deployed new FTS3 version on fts3.cern.ch (already running at RAL) fixing several bugs including e.g. checksumming.
- ATLAS: after August upgrades, RAL FTS3 server stable, going to put back Tier-1 transfers on this server after checking with developers
- CMS: testing new FTS3 server at IN2P3 in Debug, increasing Debug load on CERN FTS3 server.
- LHCb: using FTS3 instance at CERN after the bugfix with the timestamps in transfer status output.

## Tracking tools

Marten announces that the end of the year target for the decommissioning of Savannah is unlikely to be met due to other more urgent duties, but it was not a hard deadline anyway. Many trackers still need migrating.

Andrea V. adds that PH-SFT has very little manpower left on Savannah, it's already on best effort. These matters will be followed up offline.

## IPv6

- Motivation: the exhaustion of the IPv4 address space is starting to create problems to some sites (in particular CERN) and WLCG needs to have a strategy to become IPv6-ready on a timescale that fits with the needs of the sites and the experiments.
- An IPv6 validation and deployment task force for IPv6 is being formed, to work in collaboration with the HEPiX IPv6 working group on these aspects:
  - ◆ Define realistic IPv6 deployment scenarios for experiments and sites (in progress)
  - ◆ Maintain a complete list of clients, experiment services and middleware used by the LHC experiments and WLCG (in progress)
  - ◆ Identify contacts for each of the above and form a team of people to run tests
  - ◆ Define readiness criteria and coordinate testing according to the most relevant use cases
  - ◆ Recommend viable deployment scenarios
- The task force should include people active or interested in IPv6 testing, from the Tier-0, Tier-1's and Tier-2's with a sufficient variety of computing and storage technologies, middleware developers and experiment software experts. Many sites and experiments already participate to the HEPiX WG.

## XrootD Deployment

dCache xrootd door monitoring plugin (from I. Vukotic)

- the first version of the third party plugin for the XRootD detailed monitoring of dCache sites has been released in CERN WLCG repo <http://linuxsoft.cern.ch/wlcg/>
  - ◆ dcache-plugin-xrootd-monitor-5.0.0.0-0.noarch
- This plugin is mandatory to enable the detailed monitoring stream from the dCache sites having joined the XRootD federations (AAA, FAX).
- NB: this version is suitable for dCache versions v. >= 2.4
  - ◆ For dCache versions v. <= 2.2 the installation of xrootd4j-backport plugin is additionally needed (not distributed in the current RPM)
    - ◇ dedicated RPM including xrootd4j-backport will be made soon available in the WLCG repo for dCache versions v. <= 2.2

## gLExec

- 45 tickets closed and verified (6 done since last meeting), 49 still open, most on hold until mid/late autumn in line with SL6 migration
- Deployment tracking page

## AOB

The next meetings will be on October 3rd and 24th.

## Action List

1. 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.
2. Investigate how to separate Disk and Tape services in GOCDB
  - ◆ proposal submitted via [GGUS:93966](#)
3. Agree with IT-PES on a clear timeline to migrate OPS and the LHC VOs from VOMRS to VOMS-Admin
  - ◆ **in progress**
4. Experiments interested in using WebDAV should contact ATLAS to organise a common discussion
  - ◆ **input welcome by the next GDB**
5. Contact the storage system developers to find out which are the default/recommended ports for WebDAV
  - ◆ **in progress** Recommendations for WebDAV deployment

-- AndreaSciaba - 17-Sep-2013

---

This topic: LCG > WLCGopsM nutes130919

Topic revision: r27 - 2013-09-23 - MartenLitmath



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback