

# Table of Contents

<b>WLCG Operations Coordination Minutes - April 3rd, 2014.....</b>	<b>1</b>
Agenda.....	1
Attendance.....	1
News.....	1
Middleware news and baseline versions.....	1
Tier-0 and Tier-1 Grid services.....	1
Storage deployment.....	1
FTS deployment.....	3
LFC deployment.....	3
Oracle deployment.....	3
T0 news.....	4
Other site news.....	5
Data management provider news.....	5
Experiments operations review and Plans.....	5
ALICE.....	5
ATLAS.....	5
CMS.....	6
LHCb.....	6
Ongoing Task Forces and Working Groups.....	7
Tracking tools evolution TF.....	7
FTS3 Deployment TF.....	7
gLExec deployment TF.....	7
Machine/Job Features.....	7
Middleware readiness WG.....	8
Multicore deployment.....	8
perfSONAR deployment TF.....	8
SHA-2 Migration TF.....	9
WMS decommissioning TF.....	9
xrootd deployment TF.....	9
IPv6 validation and deployment TF.....	9
HTTP proxy discovery TF.....	9
Action list.....	9
AOB.....	9

# WLCG Operations Coordination Minutes - April 3rd, 2014

## Agenda

- <https://indico.cern.ch/event/309124/>

## Attendance

- Maria Alandes (chair), Nicolo Magini (secretary)
- local: Andrea Sciaba, Stefan Roiser/LHCb, Alessandro Di Girolamo/ATLAS, Michail Salichos, Alberto Aimar, Oliver Keeble, Maarten Litmaath/ALICE, Simone Campana, Hassen Riahi
- remote: Alessandro Cavalli/CNAF, Christoph Wissing/CMS, Di Qing/TRIUMF, Jeremy Coles/GridPP, Maite Barroso/Tier-0, Rob Quick/OSG, Valery Mitsyn/JINR, Yury Lazin, Alessandra Forti, Antonio Perez/PIC, Burt Holzman/FNAL

## News

- The next WLCG Collaboration Workshop will take place from the 7th to the 9th of July in Barcelona. Some details are already available in the Indico Agenda. Registration will open next week. More details will be announced at the GDB next Wednesday. As far as the agenda is concerned, we are working on it and we will contact experiments and the relevant people in the upcoming weeks.
- The next WLCG Operations Planning meeting will take place on the 17.04.2014. Experiments should expose their future plans for the next three months.

## Middleware news and baseline versions

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

- Baseline version of gliteWMS (still used by SAM) to be checked and updated after the meeting.

## Tier-0 and Tier-1 Grid services

### Storage deployment

Site	Status	Recent changes	Planned changes
CERN	<b>CASTOR:</b> v2.1.14-11 and SRM-2.11-2 on ATLAS, ALICE, CMS and LHCb <b>EOS:</b> ALICE (EOS 0.3.4 / xrootd 3.3.4) ATLAS (EOS 0.3.8 / xrootd 3.3.4 / BeStMan2-2.3.0) CMS (EOS 0.3.7 / xrootd 3.3.4 / BeStMan2-2.3.0) LHCb (EOS 0.3.3 / xrootd 3.3.4 / BeStMan2-2.3.0 (OSG pre-release))	CASTOR LHCb upgraded.	
ASGC		None	None

	CASTOR 2.1.13-9 CASTOR SRM 2.11-2 DPM 1.8.7-3 xrootd 3.3.4-1		
BNL	dCache 2.6.18 (Chimera, Postgres 9.3 w/ hot backup) http (aria2c) and xrootd/Scalla on each pool	None	None
CNAF	StoRM 1.11.3 emi3 (ATLAS, LHCb, CMS)		
FNAL	dCache 2.2 (Chimera, postgres 9) for disk instance; dCache 1.9.5-23 (PNFS, postgres 8 with backup, distributed SRM) for tape instance; httpd=2.2.3 Scalla xrootd 3.3.6-1 EOS 0.3.21-1/xrootd 3.3.6-1.sl5 with Bestman 2.3.0.16		Planning to upgrade dCache tape instance to 2.2+Chimera soon
IN2P3	dCache 2.6.18-1 (Chimera) on SL6 core servers and pool nodes Postgres 9.2 xrootd 3.3.4 (Alice T1), xrootd 3.3.4 (Alice T2)		
JINR-T1	dCache <ul style="list-style-type: none"> <li>• srm-cms.jinr-t1.ru: 2.6.24</li> <li>• srm-cms-mss.jinr-t1.ru: 2.2.24 with Enstore</li> </ul> xrootd federation host for CMS: 3.3.6		
KISTI	xrootd v3.3.4 on SL6 (redirector only; servers are still 3.2.6 on SL5 to be upgraded) for disk pools (ALICE T1) xrootd 20100510-1509_dbg on SL6 for tape pool xrootd v3.2.6 on SL5 for disk pools (ALICE T2) dpm 1.8.7-4	xrootd v3.3.4 installation on SL6	xrootd upgrade planned to 3.3.4 on SL6 for all xrootd nodes (including tape pool; waiting for procurement of a SAN switch)
KIT	dCache <ul style="list-style-type: none"> <li>• atlasrm-fzk.gridka.de: 2.6.21-1</li> <li>• cmssrm-kit.gridka.de: 2.6.17-1</li> <li>• lhcb-srm-kit.gridka.de: 2.6.17-1</li> </ul> xrootd <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 redirector 3.3.3-1</li> </ul>		
NDGF	dCache 2.3 (Chimera) on core servers. Mix of 2.3 and 2.2 versions on pool nodes.		
NL-T1	dCache 2.2.17 (Chimera) (SURFsara), DPM 1.8.7-3 (NIKHEF)		
PIC		None	None

	dCache head nodes (Chimera) and doors at 2.2.23-1 xrootd door to VO servers (3.3.4)		
RAL	CASTOR 2.1.13-9 2.1.14-5 (tape servers) SRM 2.11-1		
RRC-KI-T1	dCache 2.2.24 + Enstore (ATLAS) dCache 2.6.22 (LHCb) xrootd - EOS 0.3.19 (Alice)		
TRIUMF	dCache 2.6.21		

## 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	None	None
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		
JINR-T1	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	1.8.3.1-1 for T1 and US T2s	SL6, gLite	ORACLE 11gR2	ATLAS	None
CERN	1.8.7-3	SLC6, EPEL	Oracle 11	ATLAS, LHCb, OPS, ATLAS Xroot federations	

## Oracle deployment

- Note: only Oracle instances with a direct impact on offline computing activities of LHC experiments are tracked here
- Note: an explicit entry for specific instances is needed only during upgrades, listing affected services. Otherwise sites may list a single entry.

Site	Instances	Current Version	WLCG services	Upgrade plans
CERN	CMSR	11.2.0.4	CMS computing services	Done on Feb 27th
CERN	CASTOR Nameserver	11.2.0.4	CASTOR for LHC experiments	Done on Mar 04th
CERN	CASTOR Public	11.2.0.4	CASTOR for LHC experiments	Done on Mar 06th

CERN	CASTOR Alicestg, Atlasstg, Cmsstg, LHCbstg	11.2.0.4	CASTOR for LHC experiments	Done: 10-14-25th March
CERN	LCGR	11.2.0.4	All other grid services (including e.g. Dashboard, FTS)	Done: 18th March
CERN	LHCBR	12.1.0.1	LHCb LFC, LHCb Dirac bookkeeping	Done: 24th of March
CERN	ATLR, ADCR	11.2.0.4	ATLAS conditions, ATLAS computing services	Done: April 1st
CERN	HR DB	11.2.0.3	VOMRS	upgrade to 11.2.0.4 planned for 14th April
CERN	CMSNR_ADG	11.2.0.3	CMS conditions (through Frontier)	TBA: upgrade to 11.2.0.4 (tentatively May)
BNL		11.2.0.3	ATLAS LFC, ATLAS conditions	TBA: upgrade to 11.2.0.4 (tentatively June)
RAL		11.2.0.3	ATLAS conditions	TBA: upgrade to 11.2.0.4 (tentatively June)
IN2P3		11.2.0.3	ATLAS conditions	TBA: upgrade to 11.2.0.4 (tentatively April)
TRIUMF	TRAC	11.2.0.4	ATLAS conditions	Done

- LFC for LHCb working fine on Oracle12

## T0 news

- A series of meetings are being held end of 2013 and during 2014 with IT experts and the 4 experiments in order to measure batch jobs efficiency depending on the location in which they are executed: Meyrin or Wigner. The conclusions are being drafted and will be discussed at next ISM (IT Services Meeting), on the 14th of April. After that they will be reported elsewhere. Stay tuned.
  - ◆ Some derived actions already in place: Started to monitor systematically job efficiencies on lxbatch, with automated notification.
- The WMS service at CERN is being decommissioned according to the announced timeline. The CMS and shared instances are being drained this week. The SAM instance will be kept till the replacing mechanism is in place, estimated by end of June.
- The batch capacity migrated to SLC6 is 65%, expecting 5% more at the end of next week
- Batch and plus upgraded to CVMFS client 2.1 in line with the 1st of April deadline.
- All mw services up to date according to WLCG baseline.
- The central JIRA instance is currently running JIRA 6.1.7. and hosts more than 230 projects. For Savannah project migration, we recommend Open Source projects (notably for Grid software) to go to the "Open" instance, and other projects to the central instance. (The PH/SFT Savannah team is handling the migration.)
- After some weeks of testing VOMS admin, and after finding some bugs to be fixed and some features that need further understanding/changing (because of their different behaviour to VOMRS), it has

been agreed with the developers that we need a new voms-admin release before going to production.

- Discussion:
  - ◆ Maite reminds that the schedule for the migration to SLC6 is coupled to the migration to the Agile Infrastructure, realistically ~5%/week can be done, goal is to reach 100% as soon as possible. When ~90% is reached, discussions with experiments will be needed to understand which dedicated resources will still need to stay on SLC5.
  - ◆ Maite says that the voms-admin developers have not yet provided a date for the new release, it will be discussed at the next meeting.

## Other site news

## Data management provider news

## Experiments operations review and Plans

### ALICE

- Steady activities in preparation for Quark Matter 2014 [↗](#) (May 19-24, GSI Darmstadt)
- KIT
  - ◆ around 14:00 CEST on Apr 1 the number of running jobs in MonALISA started going down from 5k+ to ~1k for unknown reasons; since then the numbers have fluctuated wildly around that low level, while the batch system typically is seeing ~3k jobs running at any time; no changes were done by ALICE and other sites are working OK; experts are investigating
- CERN
  - ◆ SLC6 job efficiencies
    - ◇ Geneva jobs looked better in the last few days, but got worse again starting around 03:00 today; the cause of these changes is unknown
    - ◇ next meeting tomorrow
- On Maite's question, Maarten answers that details about ALICE job efficiencies will be provided at the next T0 job efficiency meeting tomorrow.

### ATLAS

- MC production and analysis: activity as usual
- ADCR DB upgrade went smoothly.
- Rucio migration: now migration from LFC to Rucio as file catalog. Migration is proceeding, more than half of the ATLAS sites have been migrated. No show stoppers seen up to now.
- Rucio commissioning: we started just in the last days the commissioning of the various Rucio services. Injection of the data (Automatix), creation of a subscription and it's translation into rules for each file (transmogrieffier), rule evaluation (judge), data transfer (conveyor through FTS3).
- DataTransfer issues: observed few links with "slow transfers" (order of 0.5MB/s). we can share our understanding presented at one ADC weekly meeting few days ago. we are trying to understand within ATLAS and PerfSONAR expert how to proceed, identify, track, follow up these issues, we believe this could be a WLCG activity.
- CERN low efficiency jobs: under investigation. 2 issues under the spotlight: 1) ATLAS pilot jobs are(were) not properly cleaning up after the main process finish (successfully), there were leftover processes (becoming like daemons) which were keeping the batch system slot busy. 2) CREAM report back to ATLAS that the job is successfully finished while LSF reports it still running. This is under investigation by Ulrich et al.
- CVMFS cache: we observed failures at CERN due to the fact that one of the ATLAS file is 2.2GB and the default shared cache was set to 2GB. Slightly different issue, but related: it seems that with

CVMFS <2.1.17 if the cache is not smaller respect the partition of the possible biggest file to be transferred failures can be observed due to partition full. 2.1.17 should have fixed this.

- FTS3: now using both CERN FTS3 and RAL FTS3 servers, 80%/20%.
- Discussion:
  - ◆ Rucio performance improved since the DB hardware migration.
  - ◆ Maria asks how many links are affected by slow transfers. Alessandro gives as example 3 UK sites to TRIUMF and BNL, e.g. Cambridge-->BNL went from good rates to bad to good again over the last months. Other experiments are not really monitoring this systematically right now.
  - ◆ Oliver asks about the current procedure to involve experts in network problems; Alessandro answers that it is in place for LHCOPN but not (yet) for LHCONE. Usually the experts at the Tier-1 on one end of the link are involved to help the Tier-2, but there is no general procedure.
  - ◆ Discussion on how to handle debugging of slow transfers to be taken offline.
  - ◆ Stefan explains that CVMFS needs a cache size at least twice as big as the size of the biggest file, this is not fixed in 2.1.17. Recommended cache size is 10 GB, Maite will follow up with Steve Traylen to check what is applied now. Reference is GGUS:102824 [↗](#) Alessandro confirms that ATLAS doesn't see failures anymore.

## CMS

- DBS2 will be switched off April 7th
- FTS2 decommissioning at CERN in August 2014 ok for CMS
- FTS3: Sites should switch their Phedex Debug agents to use RAL FTS3
- small production and processing overview
  - ◆ Heavy Ion rereco pass launched last week,
    - ◇ Currently using HLT resources at CERN, working to expand it further
    - ◇ Quite memory (3-4GB) and cpu (96h jobs) demanding, playing with job splitting to mitigate
  - ◆ Phase II upgrade MC
  - ◆ soon 13 TeV MC digi/reco
- CVMFS switch at CERN: Monday, April 14th

## LHCb

- Many thanks to all the sites who have installed already the pledged disk capacities
- Incremental stripping campaign almost finished, some remaining files to be processed at NL-T1, GRIDKA and IN2P3
  - ◆ Many thanks to all T1 sites !!!!
- Future VOMS2 server added to the VO card
- PES informed LHCb about possible LSF jobs with low efficiency. An example job has been investigated with the information available from LHCb/Dirac and no inefficiency of the payloads within this LSF job have been found. Investigations ongoing....
- Discussion:
  - ◆ Stefan would like to know about the period used for the calculation in the notifications about low efficiency jobs. In contact with Ulrich about it.
  - ◆ Alessandro mentions a feature request to dump the information in AFS to be able to review it later without digging in e-mails.
  - ◆ Need to understand who should be notified, not appropriate to send everything to the owner of the account on which the pilots run
  - ◆ Maarten comments that ALICE sees strange things in LSF like jobs not killed after days/weeks.

## Ongoing Task Forces and Working Groups

### Tracking tools evolution TF

- New release of GGUS on 26 March
  - ◆ New features: multiple site notification, CMS specific SU and forms
  - ◆ Initial test of the alarm notification failed. Repeated again on 31 March, and then it worked
- 'GGUS Shopping list' tracker migrated from Savannah to JIRA on 1st April
  - ◆ All users/tickets imported.
  - ◆ One field ('Planned Release', aka 'Fix Version') was not imported properly. Still under investigation

### FTS3 Deployment TF

- New version deployed on pilot, containing among others:
  - ◆ <https://svnweb.cern.ch/trac/fts3/wiki/ReleaseNotes#Releasecandidate>
  - ◆ Features requested in TF
    - ◇ bandwidth restrictions
    - ◇ transfer multi-hop
    - ◇ stage-in from tape only
    - ◇ activity VO shares
    - ◇ bulk cancellations via a file
  - ◆ Fixes for issues seen in CMS ASO scale testing
    - ◇ delegation issues related to voms roles, server-side timeouts with legacy client
  - ◆ Improvements to optimizer plots (increasing time range up to 7 days)
- CMS proceeding with migration to FTS3
- Michail comments that deployment to production of the new version will happen in one or two weeks if no issues are seen.

### gLExec deployment TF

- 79 tickets closed and verified, 16 still open (no change)
  - ◆ progress with a few cases
- Deployment tracking page
- Maria asks if the TF should stay open until gLExec is deployed on all sites, Maarten answers that it should be discussed in the MB

### Machine/Job Features

- Meeting last week, detailed plan for bare metal, cloud, client and bi-directional developments has been discussed and agreed within the TF
- SAM probe for checking the availability and course correctness of features has been deployed in pre-production (for LHCb)
- Deployment of the mjf client has been done in the LCG/AA area
- Discussion:
  - ◆ Maria asks about other sites which have deployed the client. Stefan explains that there is a working implementation for all batch systems except SLURM. KIT has deployed the client, NIKHEF next.
  - ◆ SLURM is deployed at ~10 sites (nordic and CSCS) but growing, need to ask them for implementation



- ◆ Stefan explains that the SAM probe is generic for all VOs, except for a minor detail (client is taken from CVMFS or web service if it is not found on the WN)

## Middleware readiness WG

- The minutes of the 3rd meeting [☞](#) held on March 18 are now available on indico direct link to the minutes [here](#) [☞](#). MariaD will make the list of **confirmed** *Volunteer* sites for the April 15 WLCG MB, relying on the Experiment wikis linked from column "Documentation pages" of the Experiment Workflows' table, provided they are also used by the Product Teams for testing, as we agreed they should be close to the MW providers AND the VOs. Candidate (but are they **confirmed**?) *Volunteer* Sites so far:
  - ◆ ATLAS: Triumph, Edinburgh, QMUL, OSG, CERN, INFN-T1.
  - ◆ CMS: Grif, CERN, Legnaro.
  - ◆ ALICE: not yet defined. We have an action on Maarten from the 18 March minutes
  - ◆ LHCb: we have no site mentioned, although their doc [☞](#) is quite detailed.
- Please contribute to the WG e-group [wlcg-ops-coord-wg-middleware@cern.ch](mailto:wlcg-ops-coord-wg-middleware@cern.ch) comments for the the Procedure Guidelines for VOs & Sites, **in particular** for the WLCG MW Officer's tasks, that will be re-discussed at the next (4th) meeting on **Thu May 15 at 10:30 CEST** in room 513-R-068 with audioconf.

## Multicore deployment

- News from experiments concerning multicore:
  - ◆ CMS: Multicore project discussed during the last CMS Comp & Offline week. Work is ongoing for the multithreaded application, therefore CMS will test its multicore submission pilots with single core jobs. Tier1 representatives have been contacted to start these tests immediately with a two step plan: functional tests, followed by real scale submission.
  - ◆ ATLAS: Atlas can already pass arguments to the batch system, but it requires to create a new panda queue for each set of parameters one may want. ATLAS is now working on making the queue creation dynamic based on the site characteristics. The current problem though is that most CREAM sites blah scripts don't pass the parameters to the batch system. Only SGE and ARC-CE sites have this capability. This needs to be looked into at WLCG level.
- Review of sites/batch system experiences with multicore jobs continues: done Condor at RAL, UGE at KIT, Torque/Maui at Nikhef. Next week SLURM.
- Next step: review experience in CMS and ATLAS shared sites when handling multicore jobs from both VOs.

## perfSONAR deployment TF

- Deadline for perfSONAR installation has passed (April 1st). 9 sites missing out of 111.
  - ◆ BelGrid-UCL: asked for SLC6 installation, pointed to <https://code.google.com/p/perfsonar-ps/wiki/Level1and2Install> [☞](#) RPM bundle (unsupported).
  - ◆ GR-07-UOI-HEPLAB: no hardware, on hold.
  - ◆ GoeGrid: no reply, 4 reminders
  - ◆ ICM: "We do not have free resources to deploy perfSonar", ticket closed.
  - ◆ MPPMU: procuring hardware
  - ◆ RO-11-NIPNE: site under upgrade on 09/01/2014, no news since then (2 reminders)
  - ◆ T2\_Estonia: under installation/configuration
  - ◆ TECHNION-HEP: first reply yesterday (3 reminders).
  - ◆ USCMS-FNAL-WC1: installed and configured (since long time), not publishing in OIM
- Still many services need attention (wrong configuration/old releases/firewalls).
- Future of TF should be discussed at the planning meeting.

- Maria asks why the FNAL server is not published in OIM. Burt answers that the FNAL perfSONAR deployment was consolidated recently, this was possibly overlooked. The reminders by Lucy and Marek were somehow missed, Burt will follow up.

## SHA-2 Migration TF

- the EGI Operations Portal VO cards for the experiments have been updated with the details of the future VOMS servers

## WMS decommissioning TF

- CERN WMS instances for experiments are being drained as of 13:53 CEST on April 1
  - ◆ no issue reports were received
- SAM instances have their own timeline

## xrootd deployment TF

## IPv6 validation and deployment TF

- A BeStMan SE and and OSG CE with IPv6 have been kindly set up at Nebraska for testing in CMS. They'll soon be added to the transfer tests and the Condor-G submission tests.
- The Panda Dev instances are being made dual stack. The next step is the JEDI development (JEDI is a component of the new ATLAS production system for Run 2) to progress further before trying to run jobs.

## HTTP proxy discovery TF

## Action list

1. Investigate how to separate Disk and Tape services in GOCDB
    - ◆ proposal submitted via GGUS:93966
    - ◆ **in progress** - ticket updated, current solution to be validated.
      - ◇ Some of the T1 sites are adding `SRM.nearline` entries as desired.
      - ◇ Downtime declaration tests to be done.
      - ◇ Experiences to be reported in the ticket.
  2. Agree with IT-PES on a clear timeline to migrate OPS and the LHC VOs from VOMRS to VOMS-Admin
    - ◆ **in progress**
      - ◇ voms-admin test cluster is available
      - ◇ experiment VO managers have given feedback
      - ◇ VOMS service managers are in contact with developers to get major issue(s) fixed
      - ◇ <https://twiki.cern.ch/twiki/bin/view/GridServices/VomsAdminTestInstance>
- Stop tracking item on Disk/Tape services, to be followed up on LHCb side
  - Stop tracking item on voms-admin, now included in regular Tier-0 reports
  - Maria comments that a new action on networking issues might need to be added

## AOB

-- NicoloMagini - 31 Mar 2014

---

This topic: LCG > WLCGOpsMinutes140403

Topic revision: r35 - 2018-02-28 - MaartenLitmaath



Copyright &© 2008-2021 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](#)