

Table of Contents

WLCG Operations Coordination Minutes - January 30, 2014.....	1
Agenda.....	1
Attendance.....	1
News.....	1
Middleware news and baseline versions.....	1
Tier-1 Grid services.....	1
Storage deployment.....	1
FTS deployment.....	3
LFC deployment.....	3
Other site news.....	3
SAM migration plan.....	3
Experiments operations review and Plans.....	4
ALICE.....	4
ATLAS.....	4
CMS.....	5
LHCb.....	5
Ongoing Task Forces and Working Groups.....	5
FTS3 Deployment TF.....	5
gLExec deployment TF.....	6
IPv6 validation and deployment TF.....	6
Machine/Job features TF.....	6
Middleware readiness WG.....	6
Multicore deployment.....	6
perfSONAR deployment TF.....	7
SHA-2 Migration TF.....	7
Tracking tools evolution TF.....	7
WMS decommissioning TF.....	7
xrootd deployment TF.....	8
Action list.....	8

WLCG Operations Coordination Minutes - January 30, 2014

Agenda

- <https://indico.cern.ch/conferenceDisplay.py?confId=295761>

Attendance

- Local: Simone, Markus, Raja, Felix, Oliver, Alberto, Marian, Maarten, Maite, Michail, Maria, Alessandro, Nicolo'
- Remote:

News

- Reminder about WLCG Ops F2F on February 11th, next regular meeting after 2 weeks.

Middleware news and baseline versions

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

- Highlights: WMS baseline downgraded to 3.6.1 for issues; APEL baselines added after meeting
- OpenSSL issue
 - ◆ WMS also needs new version of `glite-px-proxyrenewal`
 - ◇ new rpms to be released in next EMI-2 and -3 Updates (foreseen for next week)
 - ◆ Globus RFE opened by Simon Fayer of Imperial College
 - ◇ Globus **library** should also move to 1024-bits default!

Tier-1 Grid services

Storage deployment

Site	Status	Recent changes	Planned changes
CERN	CASTOR: v2.1.14-5 and SRM-2.11-2 on all instances EOS: 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 / BeStMan2-2.2.2)	EOS BeStMan updates for ATLAS,CMS (needed workaround for 512bit-proxy)	
ASGC	CASTOR 2.1.13-9 CASTOR SRM 2.11-2 DPM 1.8.7-3 xrootd 3.3.4-1	None	None

BNL	dCache 2.6.18 (Chimera, Postgres 10 w/ hot backup) http (aria2c) and xrootd/Scalla on each pool	None	None
CNAF	StoRM 1.11.3 emi3 (ATLAS, LHCb) StoRM 1.11.2 emi3 (CMS)	Upgraded Atlas and LHCb	
FNAL	dCache 1.9.5-23 (PNFS, postgres 8 with backup, distributed SRM) httpd=2.2.3 Scalla xrootd 2.9.7/3.2.7.slc EOS 0.3.2-4/xrootd 3.3.3-1.slc5 with Bestman 2.2.2.0.10	dCache disk pool at 2.2, currently populating metadata	Today: EOS 0.3.15/xrootd 3.3.6; in a few weeks, dCache tape pool at 2.2
IN2P3	dCache 2.6.18-1 (Chimera) on SL6 core servers and pool nodes Postgres 9.2 xrootd 3.0.4 (mixed with 3.3.4)	Upgrade of 75% of Alice T2 xroot nodes to 3.3.4	Upgrade to 3.3.4 for Alice T1 & T2
KISTI	xrootd v3.2.6 on SL5 for disk pools xrootd 20100510-1509_dbg on SL6 for tape pool dpm 1.8.7-3		
KIT	dCache <ul style="list-style-type: none"> • atlassrm-fzk.gridka.de: 2.6.17-1 • cmssrm-kit.gridka.de: 2.6.17-1 • lhcbssrm-kit.gridka.de: 2.6.17-1 xrootd <ul style="list-style-type: none"> • alice-tape-se.gridka.de 20100510-1509_dbg • alice-disk-se.gridka.de 3.2.6 • ATLAS FAX xrootd proxy 3.3.3-1 	None	None
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	dCache head nodes (Chimera) and doors at 2.2.23-1 xrootd door to VO servers (3.3.4)		On 25th/Feb. Upgrade to dCache 2.6.20
RAL	CASTOR 2.1.13-9 2.1.14-5 (tape servers) SRM 2.11-1	none	Feb/Mar: Upgrade to 2.1.14-5
TRIUMF	dCache 2.2.18		Upgrade to dCache 2.6.20 in Feb.
JINR-T1	dCache <ul style="list-style-type: none"> • srm-cms.jinr-t1.ru: 2.6.19 • srm-cms-mss.jinr-t1.ru: 2.2.23 		

	with Enstore		
	xrootd federation host for CMS: 3.3.3		

- Raja: RAL waiting for CASTOR 2.1.14 patch to fix issues with T10KD tape drives before upgrading

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		
KIT	2.2.8 - transfer-fts-3.7.12-1	None	None
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	None	None
RAL	2.2.8 - transfer-fts-3.7.12-1		
TRIUMF	2.2.8 - transfer-fts-3.7.12-1		Upgrade to dCache 2.6.20 In Feb
JINR-T1	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	The DBAs are planning to upgrade the (LHCBR) DB to Oracle 12 soon so the LFC has to be tested with this DB version and upgraded if needed beforehand.

- CERN LFC: The DB people have mentioned that they plan to upgrade the (LHCBR) DB to Oracle 12 soon so the LFC would have to be tested with this Oracle version and updated if needed before the upgrade.
- Maite: Fabrizio Furano for LFC development is in the loop.
- Simone: organize discussion with IT-DB about how upgrade plans to Oracle12 affect all WLCG services - motivation, schedule, planning

Other site news

SAM migration plan

- Marian presents plan to split SAM services for WLCG (at CERN) and EGI (at consortium)
- Comments after the presentation:
 - ◆ Services split to focus on different requirements

- ◆ SAM code will also be forked; EGI development taken over by consortium, possibly replacing Nagios eventually
- Simone: migration plan looks safe. No objection from the audience

Experiments operations review and Plans

ALICE

- gearing up for Quark Matter 2014 [↗](#) (May 19-24, GSI Darmstadt) with some special MC productions and more active analysis
- CERN
 - ◆ Wigner vs. Meyrin job failure rates and CPU/wall-time efficiencies
 - ◇ network issue suspected (Wigner jobs access EOS data)
 - ◇ CPUs at Wigner have lower HS06 ratings
 - ◇ to be continued...
- KIT
 - ◆ occasional data corruption due to silent PFN clashes when multiple Xrootd servers create files concurrently
 - ◇ PFN now based on a time stamp (at least by default), whereas it used to be a hash of the LFN
 - ◇ the Xrootd servers at KIT all see the whole name space
 - ◆ affected ~0.9% of data written since new servers were configured in Sep
 - ◇ 47,773 out of 5,303,211 files
 - ◇ ~1.05 TB
 - ◆ new servers read-only for now
 - ◆ cleanup of SE and catalog in preparation
 - ◆ various solutions being checked and compared
- Comments: Lower HS06 ratings in Wigner known and due to different hardware purchased; could explain higher rate of timeout failures. Likely cause for lower CPU efficiency is network instead.

ATLAS

- Rucio renaming campaign almost over.
- Rucio commissioning has started. We are now sending HC jobs to few volunteered sites (today one, we will submit to more in the coming days). We expect 2/4 weeks to validate the new DQ2 clients which are making transparent the introduction of the Rucio catalog instead of the LFC, then we will gradually move sites from LFC catalog to Rucio catalog. In a second moment we will move from DQ2 clients (CLI and API) to Rucio clients.
- Panda Jedi for analysis is under commissioning. Again HC jobs are sent to verify the full chain, next step is to include few power users, then we open to the (ATLAS) world. 4 weeks is the estimated time.
- Long standing issue (even if low % of errors): CERN-PROD CERN AI batch nodes have few % of failures not understood. First was a problem of space, now we observe many jobs killed due to memory limitation: this was not happening previously, it's not related to analysis or production, we need support to solve the issue, we will see which kind of way to track and solve the issue is the best.
- DC14 simulation started on 1st of January after multicore mode of execution was validated in December. A preliminary commissioning of multicore sites was performed on most of the ATLAS Tier-1 sites with several Tier-2 sites contributing. In total, about 20 sites support multicore submission, half of them being dynamically configured for transparent resource allocation.
- Simone: HC for Jedi PanDA in place? Alessandro: yes, no issues in HC. Some issues in PanDA reported and fixed.

- Maarten: if all goes well with Rucio migration, can we say that LFC is not needed anymore by ATLAS e.g. from summer? Alessandro: most likely yes (need to check for tape data).

CMS

- DBS migration has been postponed
 - ◆ New date is the week starting Feb 10th
 - ◆ Production will be drained
- Tier-1s actively used for Re-Reco workflows
- Tier-2 usual MC production and user analysis
- Scheduling of SAM tests
 - ◆ gLExec test (not yet critical) is a bit difficult for Tier-1s
 - ◇ Running with plain pilot VOMS role
 - ◇ Present scheduling policy asks higher priority of production
 - ◇ Might stay queued until time out
 - ◇ Discussions are ongoing how to improve this
 - ◆ Other tests should not cause issues
 - ◇ Send with lcgadmin or production VOMS role
- Simone: waiting to sort priority issues before turning gLExec test critical? Christoph: yes

LHCb

- GGUS statistics : 19 tickets opened in the last 2 weeks,
 - ◆ 8 for failed / aborted pilots, 2 for CVMFS problems, 2 VOMS access problems,
- 5 sites supporting LHCb are still running with CVMFS client version < 2.1 (in view of needed upgraded of Stratum 0), and will be contacted by GGUS tickets
- WMS decommissioning, the last sites with indirect submission are switched over to Direct submission now
 - ◆ The CERN WMS instances will be removed for indirect submission by April as requested by the TF (relying on other WMSes if needed)
- FTS3 : waiting for python binding for 2.7
- Issues with ARC CEs :
 - ◆ Default ARC installation does not allow setting up of job environment for VOs
 - ◆ Monitoring jobs submitted to ARC is not straightforward - jobs take time to appear as submitted in the ARC monitoring
- Oliver: python2.7 bindings for FTS3 underway

Ongoing Task Forces and Working Groups

FTS3 Deployment TF

- FTS3 task force meeting on Wednesday 29th (yesterday).
- Discussed and agreed among the ones present that presently it is ok to have VO production role people able to modify the configs, e.g. apply limits to sites or links. If sites need to set limit, please contact your experiment contacts. We are happy to re-discuss the issue in few months from now once we see how often we will need this kind of requests.
- Experiments re-started increasing the load on the RAL FTS3 instance.
- Testing optimizer and implementing monitoring to understand behaviour.

- Simone: reminder about FTS3 deployment model discussion in WLCG F2F on Feb 11th. Alessandro invites all stakeholders who are not usually joining the technical discussion in the TF meetings.

gLExec deployment TF

- 73 tickets closed and verified, 22 still open
- EMI gLExec probe (in use since SAM Update 22) crashes on sites that use the tar ball WN and do not have the Perl module `Time/HiRes.pm` installed (GGUS:98767 [↗](#))
 - ◆ installation of that dependency now looks sufficient to cure the problem
 - ◆ a proper fix is still to be decided
- Deployment tracking page

IPv6 validation and deployment TF

- Simone: results of the IPv6 F2F to be digested, report at next meeting

Machine/Job features TF

- NR

Middleware readiness WG

- Next meeting on Feb 6 at **15:30** CET: agenda [↗](#) in particular about how to involve experiments and sites. WG members from the experiments are called to contribute material for table: https://twiki.cern.ch/twiki/bin/view/LCG/MiddlewareReadinessArchive#Experiment_workflows. Namely:
 1. which experiment application can be adapted to participate in middleware testing
 2. which middleware product will be tested by that
 3. sites willing to play the game
 4. which of the experiment's (WLCG-accessible please!) twiki or other doc. page will contain the process
 5. additional comment, if any
- Middleware Readiness home page is evolving:
 - ◆ Product Teams info
 - ◆ Experiment workflows
- Maria on point 4: docs on experiment validation tests are maintained by exps but should be public to be useful for the community

Multicore deployment

- Meetings started 2 weeks ago. There will be one every week on Tuesday at 2:30 pm. First few weeks already scheduled in indico
- First meeting we gave a general overview of the problem and discussed how to proceed
 - ◆ Experiments more specific requirements. ATLAS was this week, CMS will be next week
 - ◆ Sites batch system scheduling mini workshops
 - ◆ Common definitions to avoid ambiguity
- Good participation and quite lively discussions so far.
- October 2014 proposed by TF coordinators as a target date for a functional system to be deployed, which is agreed by the experiments. A more refined list of milestones should be discussed in the coming meetings.
- Summary will be given at the February pre-GDB WLCG Ops Coord F2F meeting.

perfSONAR deployment TF

- Next release of perfSONAR-PS should be out next week (v3.3.2). Lots of minor fixes and improvements. All sites should update to this once it is ready
- MaDDash testing ongoing at <http://maddash.aglt2.org/maddash-webui> (Note: "Dashboard" is a menu where you can select various meshes). Graphs are now working.
- OMD testing ongoing at <http://maddash.aglt2.org/WLCGperfSONAR/> (Contact Shawn if you need login credentials)
- Discussion in progress on what services we want OSG to setup for WLCG
- Sites that are not deployed, not upgraded or not using the mesh-configuration are getting tickets

SHA-2 Migration TF

- EOS SRM for LHCb **not** yet OK
 - ◆ waiting for a patch to let the new SRM support the "root" protocol expected by LHCb jobs
 - ◇ OSG BeStMan support ticket opened
- `voms-proxy-init` on lxplus crashes on creating SHA-2 **RFC** proxies
 - ◆ discovered by CMS
 - ◆ works OK with Java-based version provided by `voms-clients3`
- VOMRS
 - ◆ the VOMS-Admin test cluster is **not** yet available!
 - ◇ despite what was reported here earlier and discussed during the meeting... 🤔
 - ◆ the host certs of the future VOMS service are from the new SHA-2 CERN CA
 - ◇ we will do a campaign to get the new servers recognized in `LSC` files across the grid
 - ◇ for convenience we probably will provide such files also in `rpms`
- Discussion on proxies:
 - ◆ Maarten: RFC migration was on hold one year ago, but eventually would like to migrate from old legacy proxies
 - ◆ Raja: won't `voms3` client bring issues with java version? Maarten: maybe ask devs to revive C++ client
- Discussion on `voms-admin`:
 - ◆ MariaD, Markus: Missing integration with HR DB is serious, priority should be on integration with HR DB
 - ◆ T0 no longer present at the meeting to comment on status of the integration
 - ◆ In parallel, VO managers can start to get familiar with the new interface when the `voms-admin` test instance is available

Tracking tools evolution TF

WMS decommissioning TF

- Deadline: end of April to decommission CMS and shared instances
- usage of the CMS WMS at CERN has remained low
- VOs using the "shared" WMS instances have been contacted for migration plans
 - ◆ LHCb and ILC reported no concerns about the timeline
- Maarten reminds about recent incident with CMS WMS
- Raja confirms that timeline is OK for LHCb

xrootd deployment TF

Action list

1. Investigate how to separate Disk and Tape services in GOCDB
 - ◆ proposal submitted via GGUS:93966 [↗](#)
 - ◆ **in progress** - ticket updated
 2. Agree with IT-PES on a clear timeline to migrate OPS and the LHC VOs from VOMRS to VOMS-Admin
 - ◆ **in progress**
- Comment on Disk/tape separation in GOCDB: Raja - it seems that it worked for SARA earlier (tape downtime didn't cause banning, need to double-check)
 - Simone: expand voms-admin action item

-- SimoneCampana - 17 Jan 2014

This topic: LCG > WLCGOpsMinutes140130

Topic revision: r32 - 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