

Table of Contents

WLCG Operations Coordination Minutes - 6th December 2012.....	1
Agenda.....	1
Attendance.....	1
Middleware news and baseline versions (Nicolò).....	1
Task Force reports.....	1
CVMFS (Stefan).....	1
SHA-2 migration (Maarten).....	2
Middleware deployment (Maarten).....	3
FTS 3 integration and deployment (Nicolò).....	4
Squid monitoring (Dave Dykstra).....	4
News from other WLCG working groups.....	4
Tracking tools (MariaD).....	4
Experiment operations review and plans.....	5
ALICE (Maarten).....	5
ATLAS (Guido).....	5
CMS (Oliver).....	6
announcements.....	6
developments.....	6
issues.....	6
questions.....	7
LHCb (Stefan).....	7
GGUS tickets.....	8
Tier-1 Grid services.....	8
Storage deployment.....	8
FTS deployment.....	9
LFC deployment.....	9
Other site news.....	9
Data management provider news.....	9
AOB.....	9
Action list.....	9

WLCG Operations Coordination Minutes - 6th December 2012

Agenda

- <http://indico.cern.ch/conferenceDisplay.py?confId=219493>

Attendance

- *Local: Maria Dimou (chair), Andrea Valassi (secretary), Oliver Gutsche, Michail Salichos, Maarten Litmaath, Nicolo Magini, Maite Barroso, Stefan Roiser, Ikuo Ueda, Guido Negri, Alessandro Di Girolamo*
- *Remote: Ian Collier, Massimo Sgaravatto, Dave Dykstra, Jeremy Coles, Gareth Smith, Andreas Petzold, Di Qing, Andrew Sansum, Christoph Wissing, Alessandro Cavalli, Michel Jouvin, Burt Holzman*

Middleware news and baseline versions (Nicolò)

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

See the Tier_1_Grid_services section below for the currently deployed versions at Tier1 sites.

Discussion about LFC

- Maite: LFC glite end of support end November, LFC instances at CERN will be upgraded after the end-of-year closure to the latest EMI2/SLC6 version. Nicolo: will this use the same Oracle? Maarten: yes, the same Oracle.
- Nicolo: question for BNL, when will you upgrade? Maria: please open a GGUS ticket on BNL so that they answer this question.

Discussion about EMI2 WN

(https://twiki.cern.ch/twiki/bin/view/LCG/WLCGBaselineVersions#IMPORTANT_NOTES_ABOUT_EMI_2_WN_U)

- Nicolo: `gsidcap` bug was fixed. Sites (especially ATLAS) should apply latest updates to WN. Exceptions are sites that need to install from tarballs.

Discussion about EMI2 UI

(https://twiki.cern.ch/twiki/bin/view/LCG/WLCGBaselineVersions#IMPORTANT_NOTES_ABOUT_EMI_2_UI)

- Nicolo: latest version of the EMI-2 UI 2.0.1 is mostly OK but it is still affected by a bug in job submission to gliteWMS, which is scheduled to be fixed in the December Update of EMI-2. The UI tarball is not available yet.

Task Force reports

CVMFS (Stefan)

- 96 sites targeted by taskforce
- 23 Sites deployed (+5 since last meeting)
 - ◆ ZA-UJ, UB-LCG2, RO-16-AIC, ...

- 36 sites missing to report on deployment plan or issues (-5 since last meeting)
 - ◆ Split by country
 - ◇ 10 Italy
 - ◇ 4 Russian
 - ◇ 3 Germany
 - ◇ 2 Croatia
 - ◇ 2 France
 - ◇ 2 Poland
 - ◇ 1 Bulgaria
 - ◇ 1 Canada
 - ◇ 1 China
 - ◇ 1 Cyprus
 - ◇ 1 Greece
 - ◇ 1 Hungary
 - ◇ 1 India
 - ◇ 1 Pakistan
 - ◇ 1 Romania
 - ◇ 1 Spain
 - ◇ 1 UK
 - ◇ 1 Ukraine
 - ◇ 1 USA (??)
- Alice is joining the task force, site names are going to be supplied and be taken into account

Discussion

- Stefan: a meeting is scheduled for next week
- Stefan: there are still 36 sites that have not filled in the twiki
<https://twiki.cern.ch/twiki/bin/view/LCG/CvmfsDeploymentStatus>.
 - ◆ Michel: you could send a reminder on the GDB mailing list, then we can contact all sites individually. MariaD: you can also directly open GGUS tickets on all sites individually, you could create one ticket and ask TPM to clone it if you provide a list of target sites. Stefan: thanks, will discuss this offline, it would be nice to have a tiered approach with regional contacts involved.
 - ◆ AlessandroC for the 10 Italian sites: will follow up with the operations people in the NGI, but we should still open tickets on all sites individually.

SHA-2 migration (Maarten)

- SHA-2 readiness verification:
 - ◆ Maarten and Steve checking VOMS/VOMRS readiness.
 - ◆ New SHA-2 CERN CA will appear in Jan.
 - ◇ But not yet in IGTF...
 - ◆ Work with EGI starting in Jan.
 - ◆ OSG are already ahead.
- We need to start verifying the **RFC-proxy** readiness for all relevant grid services and clients:
 - ◆ RFC proxy and SHA-2 signature support in WLCG middleware

Generic links:

- Sep GDB update [↗](#)
- 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>

More details from Maarten:

- Tried with Steve Traylen (CERN VOMS manager) to add to the list of supported CAs a SHA-2 CA for which Maarten has a certificate. Did some progress, but some issues are pending.
- Will have a new CERN CA supporting SHA-2 sometime in January, but this will not be automatically distributed in IGTF.
- EGI will also try to ramp up the infrastructure for SAM in January, which will be a natural way to test the readiness of the end points. OSG seems to be quite advanced.
- The good news is that we can start verifying the RFC proxies. But we have no news about dcache. We still hope that we can decouple these two things, so our best plan at the moment is to first prepare the infrastructure for RFC proxy use and then move dcache and BeStMan.

Discussion:

- Stefan: with SHA-2, will the CA name change? There were some issues with the UK CA in the past for similar reasons. Maarten: IIUC, yes the CA name will change. Stefan: will the user DN remain the same? Maarten: yes the user DN should remain the same. Anyway will check and report at the next meeting, indeed there may be problems in applications that depend on the CA/DN combination. Stefan: DN changes could definitely cause problems. Action on Maarten to report at the next meeting.

Middleware deployment (Maarten)

- Worker Node testing for WLCG

Matt Doidge provided an update on the tar ball WN this afternoon:

- We had had a reasonable sized cluster at Lancaster running the current version of the SL5 EMI2 tarball effectively in production without any problems for a few weeks now.
- The barebones of the tarball is just about ready. David Smith is currently working on smoothing the rough edges and getting the tarball build to work for SL6. My current task is trying to get yaim to automate the setup process - which is looking hopeful as the "tar" use case still exists in the yaim code.
- The exact method of distribution is still up in the air. The options are to distribute the scripts as a "tarball generation toolkit", allowing sites to keep their tarball up to date. This also mitigates any licensing issues with distributing software from EPEL.
- The second option is to distribute the tarball as before, in a tarball hosted somewhere. The traditional way.
- Option 3 is to distribute via cvmfs, but if we take this route I can't see it being ready until the end of January. Of course none of these are mutually exclusive.
- Our aim is to have something ready, probably in the form of the "toolkit", for other sites to test by the middle of next week. That leaves little to no time for testing before the Christmas break, but the idea is that it will be there for when admins come back in the New Year.

Discussion:

- Oliver: will we do the UI tarball when we are done with the WN? Maarten: yes, UI is not much extra work when the WN is done, so we'll work on UI at the end of January after WN.

FTS 3 integration and deployment (Nicolò)

- Meeting on Nov 28th, minutes in https://svnweb.cern.ch/trac/fts3/wiki/Minutes_28_11_2012
 - ◆ FTS3 demo by developers: HTTP protocol transfers
 - ◆ Asked to install FTS3 clients in a public area, to start testing new functionality
 - ◆ Discussed how to use FTS3 as fallback in case of production FTS outage: as first step, proposal to set up an FTS3 instance with auto-tuning enabled and no manual configuration to be used as fallback
- Next meeting on December 19th: main topic will be demonstration of manual configurations of VO shares and protocols on endpoints/links

Discussion:

- Stefan for LHCb: did you discuss the installation of the client in the LCG AA AFS area? Michail: good idea, will discuss this with David.

Squid monitoring (Dave Dykstra)

- See presentation on the agenda page

Discussion:

- Ian Collier: we are very unlikely to allow MRTG monitoring of our Squids at RAL Tier1, MRTG is a very old protocol. Dave: thanks, will follow this up with you offline. MariaD: please give us an update at the next meeting on December 20 (if the meeting is confirmed).

News from other WLCG working groups

Tracking tools (MariaD)

- The proposal to remind 'Notified Sites' as well as their 'Responsible Unit' is accepted and will be implemented, as it was discussed twice in this meeting and comments were only positive. Details in Savannah:131988#comment12
- The proposal to close GGUS tickets after 15 working days in status 'Waiting for reply' will not be implemented, as comments were rather against it. For GGUS tickets to the CERN site this happens de facto due to SNOW but the active ticket monitoring prevents tickets from being wrongly assigned or forgotten. Details in Savannah:133041#comment15
- The migration of the savannah-to-ggus bridge to a GGUS-only solution for CMS was discussed on 2012/12/04 in a small development circle. This technical proposal will be discussed in CMS internally and, if accepted, dates will be put against these detailed dev. items: Savannah:131565#comment14
- The future of savannah was discussed in a tracking tools TF meeting with the savannah and jira managers on 2012/12/05. Agenda and notes from this meeting on <https://indico.cern.ch/conferenceDisplay.py?confId=219255> . Action on tracking tools TF members who own savannah projects to list them and report to the TF <wlcg-ops-coord-tf-tracktools@cern.ch> (which includes the savannah and jira developers) what they wish to do with them (freeze/migrate-to-jira/other(what)).
- GGUS Release on 2012/12/12. Following the discussion (see below) ALARM tests **shall** be done. Related tickets Savannah:134078#comment8 and Savannah:134467 .

Discussion:

- Maarten: would still advise to do alarm tests as usual, there have been cases when it was "supposed to be transparent" and it was not.

Experiment operations review and plans

ALICE (Maarten)

- CNAF: Thu Nov 29 afternoon there was a lot of contention for disk I/O on the WN, in particular slowing down jobs using CVMFS. As on the previous occasion when this problem was observed, the ALICE task queue was almost empty, thereby leading to many agents preparing SW for tasks that had started on other sites in the meantime, without new tasks taking their place. It is not clear if that fully explains the problem, though. For example, LHCb reported an independent CVMFS problem on Fri Nov 30. For the time being ALICE is again using the shared SW area instead of Torrent.

ATLAS (Guido)

Reprocessing of 2012 data on-going, close to the end with some more merging jobs and re-running of failed jobs.

- avoided data on FZK Tape as input (tape performance degradation)
- suffered from "data loss" at T1s (FZK, NDGF -- disk server incidents, RAL -- power cut)
 - ◆ ATLAS finally started exercising recovery of a job output by running a single job, rather than re-running a whole task (a long-wished function in our prodsys)

Another set or processing of special stream data from tape is to be defined soon (this month).

Follow-up within ATLAS about Frontier raised last week "*Frontier: to avoid default TCP timeout in case of service down (WLCGDailyMeetingsWeek121119#Wednesday)*"

- The Frontier client is configured with a 10 second TCP timeout, and try the next Frontier server on the list quickly if the primary Frontier server is down
- i.e. better to keep the node down, when the database behind is down, rather than rebooting with a possibility the Frontier server sending a "keep alive" command potentially causing the time to fail and retry to be longer.

Points raised at the last meetings to be followed-up

- GOCDB: WLCG-ops should review the fall-back system / procedure
- FTS: affects T2 activities largely. ATLAS requests WLCG-ops to address fall-back solution
- VOMS-GGUS synchronization for /atlas/team (WLCGDailyMeetingsWeek121029#Friday)
- Need for alert when OPN switches to backup (WLCGDailyMeetingsWeek121105#Monday)
- Twiki WLCGCriticalServices to be updated

ATLAS Distributed Computing Tier-1/Tier-2/Tier-3 Jamboree (10-11 December 2012 CERN)

- <https://indico.cern.ch/conferenceDisplay.py?confId=196649>

Discussion:

- Andrew Sansum: there was no data loss at RAL, just an unavailability. AleDG: there was a problem with 177 data files at RAL, indeed a minor problem. This was not really a data loss but data corruption. Within ATLAS, when there is data corruption we need to declare this as data loss (internal ATLAS terminology) to be able to clean the catalog. Andrew: data corruption was due to power cut while transferring, hence data was only partially received.
- Maarten: FZK reported today that there was no data loss. AleDG: agree with Guido that if files are unavailable for two weeks they are "lost" because you need to decide what to do with them, e.g. to

recreate the files that were unavailable.

- Jeremy: will report on GOCDDB issues at the next meeting. Action on Jeremy.
- MariaD: VOMS-GGUS synchronization for /atlas/team is completed. The issues were due to changes in the CA/DN pair, a feature of VOMS. Steve Traylen confirmed that next time this will be handled automatically. The problem was in VOMS and not in GGUS here.
- MariaD: the issue of FTS fall back affecting T2 activities should be discussed in the FTS task force. Maarten/Nicolo: ok will discuss this there.
- MariaD: about the need for alert when OPN switches to backup, can we issue a GGUS ticket on OPN? Maarten to Guido: please follow this up with Simone in the context of the PerfSonar task force.
- MariaD: action on myself about the updates to the WLCGCriticalServices twiki. Maarten: should discuss with Maria Girone how this should be done.

CMS (Oliver)

announcements

- Improving CVMFS support
 - ◆ all software deployment team members are now watching the cvmfs savannah squad: cmscompinfrasup-cvmfs
 - ◆ questions and support requests about CVMFS should go through this squad
 - ◆ documentation is maintained here: <https://twiki.cern.ch/twiki/bin/view/CMSPublic/CernVMFS4cms> and here: <https://twiki.cern.ch/twiki/bin/view/CMSPublic/CompOpsCVMFS>, to be consolidated
- Improving savannah ticket support
 - ◆ merged categories: "Data Operations" and "Facilities Operations" into "Facilities"
 - ◆ created squad acting as catch-all place: cmscompinfrasup-comp_ops: sites are asked to reassign ticket to this squad if they need reply from central operations and don't have a specific squad they would like to communicate with
- Support for /store/himc and /store/hidata
 - ◆ All T1 sites except FNAL and all T2 sites are asked to support /store/hidata and /store/himc for production use
 - ◆ These top level directories are used to cleanly separate Heavy Ion collision files from proton-proton collision files
 - ◆ Not all sites allow for directory creation in /store, therefore this announcement
- T2s using DPM: workarounds are needed to run CMSSW
 - ◆ Reminder that workarounds have to be re-applied and updated after upgrading middleware or OS
 - ◆ Documentation is kept up to date on TWiki: <https://twiki.cern.ch/twiki/bin/view/CMSPublic/CompOpsT2DPMInstructions>

developments

- ARCHIVE Castor Service Class at CERN ready: 10 TB per user
 - ◆ Documentation: <https://twiki.cern.ch/twiki/bin/view/CMSPublic/CompOpsCERNArchive>
 - ◆ users can store and recall files from tape but cannot run CMSSW jobs against the files

issues

- IN2P3 and other dCache T2 are suffering from SRM problem:
 - ◆ known problem, the jGlobus fix for long proxies (fixed in jGlobus 2, but dCache uses jGlobus1) requires frequent restarts of frozen SRM

- ◆ Currently IN2P3 uses cron job to check for responsiveness and restarts and waits for a fix from dCache developers (IN2P3 hasn't gotten any reply in a long time)
- ◆ Best is to find problematic DN and stop user refreshing the proxy in a certain way which creates longer and longer proxies (problematic as DN is not in the log files as freeze happens before the log files are used)
- Not all downtimes declared in OIM for OSG sites are propagated to Dashboard for a change in OIM feed
 - ◆ OSG sites can get improperly marked in "error" during a downtime
 - ◆ Follow up in SAV:134221 [↗](#)
- some instabilities with CERN CreamCEs:
 - ◆ CEs: HammerCloud test jobs aborting on CERN CREAM CEs ce206 and ce208 with reason "the endpoint is blacklisted", IN PROGRESS GGUS:89124 [↗](#)
 - ◆ CEs: low level job submission problem (< 5%), IN PROGRESS GGUS:88573 [↗](#)

questions

- - ◆ UI recommendation: gLite is recommended but not supported anymore
 - ◇ How do we proceed?
 - ◆ Also related to the UI: for progress with the UI migration towards EMI, we *urgently* need a relocatable "TAR distribution". This was always lowered in priority by EMI due to person power limitations. There was though an effort started by CERN IT and GridPP to provide this. Can we have an update on the status of a tar distribution?

Discussion:

- Nicolò: the SRM issue has been going on for months by now and is getting worse, some sites need to restart their frozen SRM once per day.
- Oliver: the question about UI's has been answered in the previous reports.

LHCb (Stefan)

Very quiet in terms of operations since last meeting.

- 2012 data reprocessing
 - ◆ has finished its currently available set of files in terms of processing
 - ◆ as of Monday the last but one set of files is going to be submitted, will keep LHCb sites (T1 + T2s) busy until Xmas break
- Simulation activities have ramped up during the time when reprocessing was going down
- Issues
 - ◆ One longer standing issue with file transfers via FTS to GRIDKA disk (and partially tape), currently investigated by experts (see also GGUS:88906 [↗](#))
 - ◆ Problem with evaluation of a site which just moved to slc6 b/c of missing dependencies on the OS level. Mostly because of specific dependencies needed for LHCb applications. Is the "meta rpm" still in use?

Discussion:

- Stefan: meta rpm is the HEPOSlibs package. Oliver/Maarten: this could be useful also for CMS and ALICE (and ATLAS). AndreaV: action on myself, will follow up.

GGUS tickets

Tier-1 Grid services

Storage deployment

Site	Status	Recent changes	Planned changes
CERN	CASTOR 2.1.13-6.1; SRM-2.11 for all instances. EOS 0.2.21/xrootd-3.2.5/BeStMan2-2.2.2 for all instances except ALICE (0.2.20)		
ASGC	CASTOR 2.1.11-9 SRM 2.11-0 DPM 1.8.2-5	None	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)	Updated the Atlas endpoints on Nov 29 it has been rolled back soon the same day because of the discover of a bug (under investigation)	
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		Site downtime from 10th to 12th of December 2012
KIT	dCache atlassrm-fzk.gridka.de: 1.9.12-11 (Chimera) cmssrm-fzk.gridka.de: 1.9.12-17 (Chimera) gridka-dcache.fzk.de: 1.9.12-17 (PNFS) 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)		
PIC	dCache 1.9.12-20 (Chimera)		
RAL	CASTOR 2.1.12-10 2.1.12-10 (tape servers) SRM 2.11-1	Upgraded remaining instances to 2.1.12	None
TRIUMF	dCache 1.9.12-19 with Chimera namespace		

FTS deployment

Site	Version	Recent changes	Planned changes
CERN	2.2.8 - transfer-fts-3.7.12-1	updated to 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		
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	1.8.2-0	SLC5, gLite	Oracle	ATLAS, LHCb	Plans to upgrade to EMI2 SLC6 after the annual closure, due to the gLite 3.2 end of support, see http://glite.cern.ch/R3.2/
CERN, test	lfc-server-oracle-1.8.3.2-1	SLC6, EMI2	Oracle	ATLAS Xroot federations	

Other site news

Nothing to add

Data management provider news

Nothing to add

AOB

- **Ongoing Vidyo problems.** At 16:00 sharp (during Maarten's SHA-2 report) and again at 16:03 we were disconnected and had to manually reconnect. Michel said that this had happened also during the GDB and was reported to Thomas Baron, who said that there is a bug in the connection of the H323 room (so only the meeting room is disconnected). Action on MariaD who will follow this up.

Action list

- *Maarten will look into CA/DN changes with SHA-2.*
- *MariaD will follow up on the ongoing Vidyo problems. Action completed, I hope you can see the internals of <https://cern.service-now.com/service-portal/view-incident.do?n=INC210753>*

- *MariaD will follow up with PES about the VOMS-GGUS synchronisation problem. This action is **Done**. Answer is: VO members changing their DN/CA pair, so far, need to update also their own Groups/Roles. Steve (the VOMS manager) offered to automate this update in future cases where a CA DN changes, thus affecting many VO members.*
 - *Jeremy will follow up the review of the fall-back procedure for GOCDB (as discussed in the ATLAS report).*
 - *MariaD will update the WLCGCriticalServices twiki.*
 - *AndreaV will follow up on the HEPOSlibs meta rpm package.*
 - *Tracking tools TF members who own savannah projects to list them and report to the TF <wlcg-ops-coord-tf-tracktools@cern.ch> (which includes the savannah and jira developers) what they wish to do with them (freeze/migrate-to-jira/other(what)).*
-

This topic: LCG > WLCGOpsMinutes121206

Topic revision: r30 - 2013-01-22 - MariaDimou



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback