

Table of Contents

WLCG Operations Planning - March 21, 2013 - minutes	1
Agenda.....	1
Attendance.....	1
Agenda items.....	1
Ongoing Task Forces Review.....	1
CVMFS (M. Cinquilli).....	1
gLExec (M. Litmaath).....	1
SHA-2 (M. Litmaath).....	1
FTS3 (N. Magini).....	2
SL6 migration (A. Forti).....	2
xrootd deployment (D. Giordano).....	3
New Task Forces: Proposals.....	3
HTTP Proxy Discovery (D. Dykstra).....	3
Cloud infrastructure testing (M. Girone).....	3
Plans and news for Tier-1 and Tier-2 sites (A. Forti).....	4
Experiments Plans.....	4
ALICE (M. Litmaath).....	4
ATLAS (S. Campana).....	4
CMS (C. Wissing).....	4
LHCb (P. Clarke).....	6

WLCG Operations Planning - March 21, 2013 - minutes

Agenda

- <http://indico.cern.ch/conferenceDisplay.py?confId=239109>

Attendance

- Local: Maria Girone (chair), Andrea Sciabà (secretary), Domenico Giordano, Christoph Wissing, Felix Lee, Maarten Litmaath, Mattia Cinquilli, Maite Barroso Lopez, Xavier Espinal, Nicolò Magini, Jakob Blomer, Luca Mascetti, Jan Iven, Massimo Lamanna. Michail Salichos
- Remote: Matt Doidge, Alessandra Forti, Peter Clarke, Malgorzata Krakowian, Frédérique Chollet, Peter Solagna, Di Qing, Dave Dykstra, Burt Holzman, Robert Frank, Gareth Smith, Massimo Sgaravatto, Rob Quick, Daniela Bauer, Michel Jouvin, Ian Fisk

Agenda items

Ongoing Task Forces Review

CVMFS (M. Cinquilli)

For details see slides.

LHCb and ATLAS set April 30 as the target date for their sites to install CVMFS: after then, software will not be installed to the old NFS shared area and jobs will be submitted only to CVMFS-enabled sites. A second deployment wave will start in spring for ALICE and CMS (no dates yet). A SAM test for CVMFS will be developed.

Christoph mentions that for CMS there is a tricky issue to solve with the Lyon site, which uses the same WNs for the Tier-1 and the Tier-2.

gLExec (M. Litmaath)

- LHCb DIRAC tests deferred until after the Easter vacation

CMS should still have a discussion in their operations meeting about the gLExec deployment.

For ALICE, it is a long term activity, most probably it will start in the second half of the year.

In ATLAS some manpower issues just arose, which will be followed up offline.

SHA-2 (M. Litmaath)

- the new CERN CA has been declared ready for a few pilot users on March 19
- next step: getting VOMS to work with it
- then: have a few more pilot users added for experiments and EGI
- more news in the coming days
- EMI/UMD compatibility table maintained by EGI:

- ◆ https://wiki.egi.eu/wiki/SHA-2_support_middleware_baseline

Concerning VOMRS and VOMS, in a discussion with Steve Traylen the preferred option seems to be to insert users manually in the DB and avoid for the time being the problems of VOMRS with SHA-2. It is not yet possible for the LHC VOs and ops to drop VOMRS now and use VOMS-Admin instead.

About the plan, after VOMS is able to deal with SHA-2 proxies, the next step is to add the new CA in IGTF and have it installed by all sites. With EGI it was agreed to use a special SAM instance to measure the infrastructure readiness and for this reason ops and the LHC VOs are a priority.

Peter asks if the ops VO is already SHA-2-ready: the problem is that due to VOMRS it is not possible to follow the usual procedure to register new certificates even if the VOMS core is SHA-2-compatible. Andrea asks why one cannot just authenticate with a normal certificate and register the DN of his new SHA-2 certificate; Maarten concurs that it is a possibility but it should be verified.

FTS3 (N. Magini)

- Stress testing
 - ◆ <https://svnweb.cern.ch/trac/fts3/wiki/StressTesting>
- Main results:
 - ◆ FTS3 with a single DB is able to sustain current global FTS2 transfers with ~20% less resources.
 - ◆ FTS3 not limited by DB or webservice but rather number of parallel url-copy processes that can be sustained by a single VM --> can continue to scale horizontally. Todo: run stress-test with fake transfers to determine DB side limitation.
- Based on these preliminary results, we see no showstopper for a single-server deployment model for all WLCG.transfers.
- Will now kickstart discussion on deployment plan in task force. Starting point for proposal is along these lines:
 - ◆ Grow fts3-pilot.cern.ch to 5 "stable" pre-production VMs, this should be able to sustain ~1/6th of the FTS2 load. Keep a corresponding number of FTS3 "development" VMs for more rapid deployment of new features.
 - ◆ Identify a corresponding fraction of sites (e.g. 2 clouds) and migrate them from FTS2 to FTS3
 - ◆ At ~monthly intervals, upgrade pilot to latest FTS3 version, add more VMs and migrate more sites.

Simone asks if the transfers via url-copy were real or simulated; Michail answers that they were real, although most failed, which does not matter for the purpose of the test. Moreover, transfer status was polled.

The stress tests were done using Oracle Express 11g with just a few cores and still the load on the database was very small. MySQL still needs to be stress tested. Maria recommends to do it as soon as possible. RAL and PIC are already using the MySQL backend.

SL6 migration (A. Forti)

For details, see the slides

Alessandra reports about the first meeting of the task force held two days ago. The main conclusions were:

1.) Sites may migrate even now, after informing their experiments
2.) Until June 1st, all sites are encouraged to test SL6
3.) After June 1st, all sites are encouraged to migrate to SL6, which gives five months to migrate the bulk of the resources before the envisaged milestone at the end of Oct

HEPOS_libs is now officially released and documented; external sites should test it if using different RHEL flavours. A proper WLCG repository should be identified, though, possibly but not necessarily at CERN, as long as it is hosted in a WLCG institution.

xrootd deployment (D. Giordano)

For details, see the slides [↗](#).

Domenico illustrates the goals of the task force, which are:

- provide support to the deployment
- coordinate the monitoring efforts
- identify common needs among experiments

Ongoing activities include:

- improving the stability of the collectors (at UCSD and CERN)
- unify the monitoring efforts (Dashboard and Data Popularity)
- improve support for xrootd monitoring in dCache and DPM
- design specific SAM tests

Ian asks if the monitoring will be in the scope of the task force. Maria answers that it will, as long as it is via tools common between the experiments. **Hence, it is agreed that xrootd monitoring is covered by the task force.**

Concerning the fact that with DPM there is no way to monitor only remote access, Domenico thinks that monitoring also local activity is a good thing as long as it can be separated at the monitoring level. For example, in EOS what is monitored for the data popularity is mostly local access. So, unless there are privacy concerns, he proposes to collect also local information.

New Task Forces: Proposals

HTTP Proxy Discovery (D. Dykstra)

For details, see indico slides [↗](#)

The mandate of the task force is: define the WLCG-wide standards for grid jobs to find out what HTTP proxy or proxies to use. The rationale is to have a single, coherent solution instead of the several, incomplete solutions that exist today.

Simone comments that there are two use cases for which this might be useful: Tier-3 sites and Clouds not in AGIS. Jakob approves and thinks it will be very useful for CVMFS, eliminating the need to configure it. He offers to join the task force.

Cloud infrastructure testing (M. Girone)

Maria brings up this topic as a followup of the discussions in January's pre-GDB and February's GDB about having a task force to coordinate testing activities on clouds. Given the several requests to deploy experiment software in clouds, Ian agrees that a task force should start ASAP, given the relative shortness of LS1.

There is some discussion about possible conflicts with the activities that report to the GDB and Michel sees a risk in having two parallel forums working on similar things and with the same people, while Ian argues that the scope should be different (focused on policies in the GDB and on actual tests - for example in the HLT farms - in the TF). It is finally agreed that Michel will draft a proposal on how to organise the work and it will

be discussed at a later time.

Plans and news for Tier-1 and Tier-2 sites (A. Forti)

- dCache 1.9.12 support by EMI was extended by four months and will end on 31-08-2013
- CVMFS > 2.1 for ATLAS by the end of April at sites that want to use the shared NFS CVMFS feature. Sites running 2.0.x versions are fine to run beyond.
- Squid upgrade for everyone by the end of April to enable the new monitoring
- xrootd requested by CMS at Tier-2's
- News from UK:
 - ◆ UK has decided to replace rfio with xrootd also for ATLAS and they are testing it independently from the FAX federated work to get practice before in a more traditional environment i.e. staging-in the input

About the CMS request to enable xrootd for file access, Christoph clarifies that for now it is a recommendation, there is no deadline set and it has nothing to do with joining an xrootd federation.

Experiments Plans

ALICE (M. Litmaath)

- Start CVMFS deployment and ramp up the usage in the course of spring

Predrag contacted IT-PES to ask them to take over the ALICE CVMFS server.

ATLAS (S. Campana)

For details, see the slides [?](#).

Simone summarises the production activities foreseen for 2013. Then he stressed the importance of monitoring; on one side the ATLAS operations and management greatly appreciate the recent Dashboard improvements but on the other side stress the need to pursue the objectives set in the Operations and Tools TEG to have a more coherent view.

Concerning SL6, the final validation is ongoing and some new requirements have been added to the VO card.

Concerning the ATLAS workload management system, JEDI will become a core component of PanDA (but this does not mean that it will become mandatory, so CMS need not worry); JEDI brings several improvements in the job description, merging, etc. (the details can be discussed offline).

In order to migrate to the new RUCIO file naming schema sites are expected to use WebDAV to rename files; the goal is to have it enabled on most sites by June and all by September. Jan points out that EOS does not allow clients to rename files, so another solution will be needed. Simone agrees and adds that there are also other sites that cannot use WebDAV.

Finally it is clarified that sites do not have to join FAX because it is still under commissioning and currently sites join only on a voluntary basis.

CMS (C. Wissing)

Short Term Plans (Weeks)

- HammerCloud
 - ◆ Running with gLite WMS and Glidein submission in parallel

Cloud infrastructure testing (M. Girone)

- ◆ Detailed comparison in the next weeks
- ◆ Switch to Glidein results for site availability calculation
- SAM Tests
 - ◆ Still use gLite WMS for submission
 - ◆ Issue with recent ARC CEs - requires EMI-3 WMS release
 - ◆ Looking into direct submission via Condor_g
 - ◇ Common submission probe with ATLAS?
- Processing on HLT farm
 - ◆ Testing is continuing and scale gets enlarged
 - ◆ Investigation of observed network bottlenecks
- Processing on Agile Infrastructure at CERN
 - ◆ Tuning submission
 - ◆ Include AI resources into real production

Medium Term Plans (Months)

- Disk/Tape separation at Tier-1 sites
 - ◆ Aim: Implementation ready by Fall 2013
 - ◆ Finalizing a commissioning program
 - ◇ Start with sites that fulfill requested functionality
- Xrootd Federations
 - ◆ Aim: Have 90% of Tier-2 ready June 1st 2013 - Fallback and included in federations
 - ◆ SAM tests for xrootd being tuned - Critical tests after June 1st
 - ◆ Redirectors should reach production quality/stability in Summer
 - ◆ Monitoring infrastructure should reach production quality in Summer
- Multicore Jobs
 - ◆ Use **existing** Multicore queues to gain production experience
 - ◆ First "Dynamic Allocation": run multiple independent single core jobs
 - ◇ Target for operation June 2013
 - ◆ Extend to "forked mode"
- SL6 migration
 - ◆ CMS is fine with current plan to move resources by Oct 2013
 - ◆ Sites are encouraged to move earlier (if there is no conflict with other VOs)
 - ◆ Move of lxplus alias in April accepted - will require some education of users (SL5 still needed for certain tasks)
 - ◆ Native SL6 CMSSW builds expected for October 2013 and production architecture will change - Requires most of the sites have moved to SL6
- Castor/EOS (C. Wissing)
 - ◆ Future Tier-0 will use exclusively EOS
 - ◆ CASTOR only used for archiving
 - ◇ Phedex subscription from EOS to CASTOR
 - ◆ No rate estimates yet
 - ◇ Expected logging rate 1kHz
 - ◇ Studies ongoing

LHCb (P. Clarke)

For details, see slides [☞](#).

LHCb's plans include 8 weeks of Incremental Stripping starting in April, which will require certain bandwidth values on the tape systems (see slides). In the medium term, by April 30 all sites must have deployed CVMFS, possibly to be used also to distribute the conditions data after LS1.

Topics currently under discussion include: tighter integration with Tier-2 sites, FTS 3 integration, federated storage (either xrootd or http), the WLCG information system (still in its early development), monitoring (planning to feed monitoring information into SAM), while keeping an eye on gLExec, SL6 and perfSONAR.

Finally, two important reviews are due to report by mid 2013: one on the distributed computing system fitness for purpose and one on the Computing Model itself.

-- AndreaSciaba - 25-Mar-2013

This topic: LCG > WLCGOpsMinutes130321

Topic revision: r18 - 2013-03-29 - AndreaSciaba



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