

Table of Contents

WLCG Operations Planning - April 17, 2014 - minutes.....	1
Agenda.....	1
Attendance.....	1
Agenda items.....	1
News.....	1
Experiments Plans.....	1
ALICE.....	1
ATLAS.....	2
CMS.....	2
LHCb.....	3
Report from WLCG Monitoring Consolidation.....	4
Ongoing Task Forces and Working Groups Review.....	4
WMS decommissioning.....	4
SHA-2.....	4
Middleware readiness.....	5
FTS3.....	5
perfSONAR.....	5
xrootd deployment.....	5
Tracking tools evolution.....	5
IPv6.....	6
Machine/Job Features.....	6
gLExec.....	6
Multicore deployment.....	6
WLCG HTTP Proxy Discovery.....	7
AOB.....	7

WLCG Operations Planning - April 17, 2014 - minutes

Agenda

- <https://indico.cern.ch/event/309125/>

Attendance

- Local: Andrea Sciabà (chair), Nicolò Magini (secretary), Maarten Litmaath (ALICE), Stefan Roiser (LHCb), Marian Babik, Felix Lee (ASGC), Vincent Brillault, Pablo Saiz, Marcin Blaszczyk, Maria Alandes, Maria Dimou, Alessandro Di Girolamo (ATLAS), Hassen Riahi, Alberto Aimar, Domenico Giordano, Markus Schulz
- Remote: Yury Lazin (RRC-KI-T1), Renaud Vernet, Shawn McKee, Vanessa Hamar (IN2P3-CC), Frederique Chollet (IN2P3-CC), Christoph Wissing (CMS), Thomas Hartmann (KIT), Burt Holzman (FNAL), Alessandro Cavalli (CNAF), Alessandra Forti, Valery Mitsyn (JINR-T1)

Agenda items

News

- WLCG workshop in Barcelona (7-9 July)
 - ◆ registration is now open (see Indico page)! Registration fee: EUR 115 (+ 40 for social dinner), deadline June 9
 - ◆ Agenda still to be defined: there is a rough draft, discussed at the last GDB, any input and suggestion is very welcome
- Task forces
 - ◆ Two task forces to be evaluated for closing: xrootd and perfSONAR. A new task force on Network monitoring?

Experiments Plans

ALICE

- KIT jobs behavior:
 - ◆ The cause of the overload of the KIT firewall and the OPN link to CERN has been found:
 - ◇ for analysis jobs the location of the WN was not propagated to the central services
 - ◇ the client then ended up using not only the local replica, but close replicas as well
 - ◆ Over the weekend a patch was developed for TAlienFile in ROOT
 - ◇ it now sends the location information explicitly
 - ◆ The patch became first available in Tuesday's analysis tag and started getting used by a few trains and users
 - ◇ the results were very good
 - ◆ The expectation is for the vast majority of analysis jobs to be using the patched code after a few days
 - ◇ this will be monitored
 - ◆ The jobs cap should then get increased again gradually over the coming days
 - ◆ Our thanks to the teams at KIT for their efforts, and to the other experiments for their patience!
- Plans for the next 3 months:

- ◆ First, continuous activities in preparation for Quark Matter 2014 (May 19-24, GSI Darmstadt).
- ◆ Then reprocessing RAW data with highest priority.
 - ◇ Starting from 2011 (p+p) and some selected periods (Pass2) from 2012.
- ◆ Accompanied by the associated MC for these periods.
 - ◇ Anchored to the new calibration resulting from the preceding RAW data pass.
- ◆ Pb-Pb reprocessing is not in the plans at this moment.
- ◆ User analysis should be less intense after QM'14.
- ◆ CERN:
 - ◇ Conclusion of SLC6 job efficiency investigations.
 - ◇ Increased use of the Agile Infrastructure.
- Alessandro comments that the KIT incident is an opportunity to review if the monitoring covers all data access (transfers, remote and local jobs) by all VOs to understand the load on a storage. Collect links to internal VO monitoring pages, check what is in the common dashboard monitoring, what can be integrated. Maarten mentions that for ALICE remote access is monitored in Monalisa. Pablo comments that Dashboard has no resources right now for local access monitoring, but it can be added to the wish list. Discussion to be continued offline on the WLCG Ops Coord monitoring mailing list.
- Alessandro asks about the metrics used by Alien to define two sites 'close' for remote processing, according to Maarten and Pablo it includes RTT measurements, domain. Includes network tests between the voboxes.

ATLAS

- Alessandro presents the status and plans of ATLAS for the next months: main points are new Tier-0, new production system, Rucio migration, database access, request to deploy FAX xrootd and HTTP/WebDAV at all sites, multicore. See slides for details.
- Alessandro clarifies that the SAM tests are not affected by the Rucio migration, because they don't interact with the catalogs, only the storage.
- Alessandro clarifies that the timeline for sites to enable multicore for ATLAS is NOW if the sites can perform dynamic provisioning. Sites who cannot do it should contact ATLAS central operations (Alessandra Forti) and discuss on a case-by-case basis.
- Alessandro explains that ~90% of the sites have joined FAX, 2 T1s (NDGF and IN2P3) and a few T2s are missing. 85% of the sites have enabled HTTP/WebDAV for the Rucio renaming, but the service is not also production-quality for remote access at all sites (e.g. EOS) since the requirements are different.
- Stefan asks the reasoning behind the request of a dedicated LSF cluster for ATLAS Tier-0, since this will introduce a partitioning of the resources. Alessandro answers that during 2012 data taking there were 10 ATLAS alarms for LSF, and the problems were not identified or were caused by other users; according to IT-PES, a smaller instance can cope better with the load. CMS is going to AI instead for the Tier-0, possibly introducing partitioning in a different way. The discussion is adjourned to the next meeting since there is no Tier-0 representative attending the current meeting.

CMS

- DBS2 has been switched off and remains off most likely
- Decommissioning of CE tags for individual CMSSW releases
 - ◆ Still to come in April
- glexec SAM test
 - ◆ Remaining scheduling issues solved
 - ◆ Want to go ahead and make it critical (beginning) of May
- CSA14
 - ◆ Tier-1 Tape Staging test in May

- ◆ Analysis Challenge in Summer (July & August)
- ◆ Preparing samples
- Multi Core
 - ◆ First production through multi core pilots at PIC
 - ◆ Technical tests ongoing at other sites
- xrootd (AAA)
 - ◆ Related SAM tests not yet critical
 - ◆ Scale testing of AAA (xrootd federation) on going concentrating on European sites
- Migration from Savannah to GGUS ongoing
 - ◆ Transfer Team and Workflow started to use it
 - ◆ Minor issues being addressed with GGUS team
 - ◆ Enlarge usage over the next weeks
- FTS3
 - ◆ Finish moving sites to FTS3 for Debug transfers
- Condor_g mode for SAM
 - ◆ SAM gLiteWMS decommissioning for end June confirmed?
- Marian confirms that the June deadline for the new Condor_g SAM probes is correct, though the current schedule is 2-3 weeks late. A prototype is available, establishing the schedule for deployment to preproduction.

LHCb

- Data Processing
 - ◆ VAC
 - ◇ Used in production with several hundred VMs at Manchester, Oxford and Lancaster
 - ◇ Infrastructure moved to CERNVM3 / SLC6
 - ◆ Unification / rewrite of the pilot framework in LHCbDIRAC
 - ◇ same infrastructure to be used at WNs, VAC, BOINC, CLOUD
 - ◇ including machine / job features
 - ◆ WMS decommissioning
 - ◇ WMS server decommissioning at CERN went without problems
 - ◇ Currently LHCb is submitting < 4 % of its pilots through WMS to small sites
 - Decommissioning of remaining sites will continue on low priority
- Data Management
 - ◆ Tier2Ds (D==Disk)
 - ◇ Many Tier2Ds are using DPM as storage technology
 - Bug of the SRM/xrootd interface was fixed earlier this year and verified for version 1.8.8 - many thanks to CBPF.br and NCBJ.pl
 - The rpm is available in EPEL / testing at http://dl.fedoraproject.org/pub/epel/testing/6/x86_64/, release in EPEL pro to follow soon
 - ◆ FTS3
 - ◇ Service is used 100 % in production since several months by LHCb
 - ◇ Client was used in "FTS2 mode" within LHCbDIRAC.
 - Planning to use the REST interface (python only, avoid Boost and other C++ dependencies)
 - ◆ File Access
 - ◇ Upcoming release of LHCbDIRAC will contain ability to use natively built xroot tURLs without going through SRM,
 - ◇ Next step will be to integrate http/webdav access, should be less work with the previous work already done but so far less endpoints available....
 - ◆ LCG file catalog to DIRAC file catalog migration
 - ◇ Migration procedure is currently being prepared, no estimate yet on the final schedule

for the migration available. The objective is still before the end of the year.

- ◆ CASTOR -> EOS migration
 - ◇ LHCb is using EOS already for production data for several months
 - ◇ Last bit missing was user data migration, which is scheduled for 22nd April - Many thanks to DSS for their support !!!!
- Infrastructure
 - ◆ IPv6
 - ◇ Test infrastructure currently being setup in LHCb
 - ◇ New LHCb representative to IPv6 WLCG TF and Hepix - Raja Nandakumar [✉](#)
 - ◆ perfSonar
 - ◇ waiting for dashboard interface to consume data
- Stefan clarifies that the timeline for sites to deploy xrootd access is when they are ready.
- Stefan confirms that the timeline for the switchover to DIRAC file catalog is before the end of the year, and LHCb intends to use it in Run2.
- Shawn comments that the dashboard in the next release of perfSonar in May will have a REST API. Data from perfSonar gathered and exposed to WLCG.

Report from WLCG Monitoring Consolidation

- Pablo presents the status of the monitoring consolidation project: recent updates, with focus on SAM3 validation and the site nagios plugin; next steps. See slides for detail.
- Feedback to be submitted to the monitoring consolidation e-group for discussion. Requests/issues on JIRA tracker.

Ongoing Task Forces and Working Groups Review

WMS decommissioning

- CERN WMS instances for experiments have been drained completely without incident
- SAM instances to be decommissioned by the end of June
 - ◆ depending on successful validation of the new job submission methods developed for SAM
 - ◇ direct CREAM submission with payload
 - ◇ Condor-G

SHA-2

- future VOMS servers campaign plans:
 - ◆ reminder broadcast to be sent around May 6 (original "deadline")
 - ◆ let SAM **preprod** instances get their proxies from the new servers to measure the readiness across WLCG
 - ◇ open tickets for failing services?
 - ◆ let SAM **production** instances use the new servers as of a hard date
 - ◇ Mon June 2?
 - ◆ send another broadcast for sites and experiments to reconfigure their UI-based services
 - ◇ remove references to the old servers
 - ◆ switch off the old servers on Tue July 1
- Andrea suggests to wait until we see how many sites are impacted before opening tickets
- Agreed to keep the TF open to track this campaign

Middleware readiness

- Following the endorsement of the first *Volunteer* sites for the MW Readiness verification effort by the 2014/04/15 WLCG MB and the approval of their rewarding method document[☞], the plan now is to **document the set-up** of these sites. Please open https://twiki.cern.ch/twiki/bin/view/LCG/MiddlewareReadinessArchive#Volunteer_Sites and contribute content for the **presence of a parallel infrastructure at the site** column.
- The draft agenda http://indico.cern.ch/e/MW-Readiness_4[☞] is available for comments. We plan to consult sys admins at the WLCG sites on the use of the MW packages' version numbers published in various places e.g. the Baseline versions' table.
- Maria Dimou explains that the "baseline version" number doesn't necessarily reflect all individual updates of the packages in the dependencies. Need to understand how to publish this.

FTS3

- FTS3 successfully managing majority of experiment transfers, agreed with all experiments on FTS2 decommissioning by August.
 - ◆ ATLAS and LHCb already 100% on FTS3, CMS to complete migration to FTS3 by June.
- Understand transfer performance with FTS3: mostly validating FTS3 and Dashboard monitoring plots to see if experiment operations have all they need to understand FTS3 transfer behavior.
 - ◆ Expected result: validate optimizer performance for bulk of transfers; spot corner cases of problematic transfers which require dedicated debugging (e.g. at network level)
- Validate submission with new clients/REST API and related new functionality. Timeline for integration of new features varies by experiment.

perfSONAR

- Shawn presents the final report of the perfSONAR task force: deployed at 205 sites by April 1st deadline, only 8 missing, but 64 still running old versions. Lessons learned and important remaining issues. See slides for details.

xrootd deployment

- Domenico presents the final report of the xrootd task force: deployment status, monitoring. Items to be followed up include keeping alive the communication between the two federations, the deployment of the dCache-XRootD monitoring plugin and the registration of the endpoints in the information systems. See slides for details.
- Andrea comments that both the perfSonar and the xrootd task forces have reached their goals. He proposes to close them and create a new task force or working group to follow up on remaining issues, increasing the scope of network monitoring beyond perfSonar and xrootd in the long term (e.g. HTTP/WebDAV, FTS3 monitoring). The working group would be the forum to establish the proper procedure to follow up network issues. Shawn and Marian to lead this activity, they are asked to propose the mandate and goals and present them at the next WLCG Ops coord meeting on May 8th. Markus suggests to call the group something like 'data access monitoring' instead of 'network monitoring' to clarify the goals.

Tracking tools evolution

- Move from savannah to JIRA done for the GGUS Shopping list tracker on 9th of April
 - ◆ Migration of experiment tracker outside scope of this TF

IPv6

- Andrea presents the status of the IPv6 task force: experiments status and plans, highlights from the HEPIX IPv6 F2F meeting. See slides for details.

Machine/Job Features

- Batch infrastructure
 - ◆ Support for ALL batch system types available, including SLURM - many thanks to NDGF)
 - ◆ Deployment plan is to test on two sites initially and then roll out to remaining sites
 - ◇ LSF: CERN (done) & second site contacted
 - ◇ Condor: USC & second site contacted
 - ◇ SGE: Gridka (done) & Imperial (done)
 - ◇ Torque/PBS: NIKHEF (done) & second site contacted
 - ◇ SLURM: script being developed
- Cloud infrastructure
 - ◆ Setting up a prototype infrastructure at CERN/Openstack (similar to what was done for CERN/LSF)
 - ◆ based on couchdb + administration tools which are currently being written
 - ◆ Later move to more / other IaaS infrastructures
- Client (mjf.py)
 - ◆ First version available at WLCG repository and LCG/AA afs area for use by sites / experiments
- Bi-directional communication
 - ◆ Currently under discussion and finalizing the structure
- see also GDB talk [↗](#)

gLExec

- 79 tickets closed and verified, 16 still open (no change)
 - ◆ slow progress with a few cases
- the current status was presented in the April 15 Management Board [↗](#)
 - ◆ presentation [↗](#)
 - ◆ the WLCG project leader proposed the task force should carry on
 - ◇ details on pages 12 and 13 of the presentation
- Deployment tracking page
- Maarten explains that the plans for the task force are to gather experience after CMS turns the gLExec test critical; then discuss with LHCb how to ramp up; follow up with other experiments.

Multicore deployment

- First review of all the batch systems completed
- Had a first phase wrap up presentation
 - <https://indico.cern.ch/event/305626/contribution/0/material/slides/0.pdf> [↗](#)
 - ◆ CMS not running multicore yet, or at least not extensively enough to asses its impact on sites
 - ◆ ATLAS still has a wavelike submission pattern which is most disruptive
 - ◆ So far the most successful model of scheduling without walltime and/or a steady stream of multicore is (dynamic) partitioning especially at sites that can - one way or another - limit the number of cores to be drained at the time.
 - ◇ FZK done it with SGE native features
 - ◇ Nikhef has done it with some creative scripting
 - ◆ Backfilling not yet possible.
 - ◆ Problem with passing parameters to batch systems in CREAM

- ◇ Nikhef has shared their blah scripts to pass parameters to maui/torque Nikhef scripts [↗](#)
 - Support is on best effort and how exotic the requests are.
- ◇ SGE works out of the box
- ◇ Most SLURM and Htcondor sites use ARC-CE

- Next steps
 - ◆ CMS/Atlas testing together
 - ◆ Trying to use other parameters like walltime at sites.
 - ◆ Second round of presentations in the new conditions

WLCG HTTP Proxy Discovery

- No report

AOB

- The next meeting on May 8th will be a regular WLCG Operations Coordination meeting.

-- NicoloMagini - 14 Apr 2014

This topic: LCG > WLCGOpsMinutes140417

Topic revision: r16 - 2018-02-28 - MaartenLitmaath



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