

Table of Contents

WLCG Operations Coordination Minutes - September 4th, 2014.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Middleware News.....	1
Oracle Deployment.....	2
Tier 0 News.....	2
Tier 1 Feedback.....	2
Tier 2 Feedback.....	2
Experiments Reports.....	2
ALICE.....	2
ATLAS.....	2
CMS.....	3
LHCb.....	3
Ongoing Task Forces and Working Groups.....	3
Network and Transfer Metrics WG.....	3
Tracking Tools Evolution TF.....	4
FTS3 Deployment TF.....	4
gLExec Deployment TF.....	5
Machine/Job Features.....	5
Middleware Readiness WG.....	5
Multicore Deployment.....	5
SHA-2 Migration TF.....	6
WMS Decommissioning TF.....	6
IPv6 Validation and Deployment TF.....	6
Squid Monitoring and HTTP Proxy Discovery TFs.....	6
Action list.....	7
AOB.....	7

WLCG Operations Coordination Minutes - September 4th, 2014

Agenda

- <https://indico.cern.ch/event/326087/>

Attendance

- local: Andrea Sciaba' (chair), Nicolo' Magini (secretary), Maria Alandes, Alberto Aimar, Alessandro Di Girolamo (ATLAS), Oliver Keeble, Michail Salichos, Andrea Manzi (MW officer), Felix Lee (ASGC), Maarten Litmaath (ALICE), Hassen Riahi, Simone Campana (ATLAS)
- remote: Andrej Filipcic (ATLAS), Michael Ernst (BNL), Stefan Roiser (LHCb), Ulf Tigersted, Alessandro Cavalli (INFN-T1), Shawn McKee, Stefano Belforte (CMS), Maite Barroso (Tier-0), Antonio Perez (PIC), Christoph Wissing (CMS), Alessandra Doria, Alessandra Forti, Andrew Lahiff (RAL), Dave Dykstra, Di Qing, Jeremy Coles, Burt Holzman (FNAL), Michel Jouvin, Josep Flix (PIC), Miguel Martinez, Yuri Lazin (RRC-KI)

Operations News

- NTR

Middleware News

- the WLCG repository [will](#) become signed in the near future
 - ◆ further news will be announced in due course
- Useful Links:
 - ◆ Baseline Versions
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines:
 - ◆ No new EMI/UMD releases in the period
 - ◆ removed FTS2 from table as decommissioned or under decommissioning everywhere
- MW Issues:
 - ◆ dCache issue with a Brazilian CA Certificate : Issue affecting LHCb in some sites PIC, Gridka, SARA, there is a workaround but a fix is also under preparation (v2.6.34)
 - ◆ Missing Key Usage extension in delegated proxy: it affects Atlas and Rucio integration , the fix for Cream UI has been postponed to October, it will be integrated in HTCondor then. This does not affect submission with RFC proxies.
 - ◆ no updates regarding the latest version of Argus, still not officially released.
- T0 and T1 services
 - ◆ FTS2 decommissioning
 - ◇ Done at TRIUMF,NL_T1, RAL
 - ◇ under decommissioning ASGC,CNAF, NDGF-T1
 - ◆ Planned changes:
 - ◇ IN2P3 : dCache upgrade to 2.6.31 on 23/09
 - ◇ NL_T1: dCache upgrade to 2.6.34 to fix issue with Brazilian certificates when available.

- Alessandro asks if there are news from OSG about how to discover HTCondor CEs in the information system. For SAM the information is taken from OIM, but for the pilot factories the question is still open. Michael Ernst will follow up with Brian Bockelman. ACTION on Alessandro to collect the information.

Oracle Deployment

- Upgrade Plans Table
- Maarten reports that the migration to GoldenGate went fine for IN2P3.

Tier 0 News

- **AFS UI:** in contact with the AFS team to get some info about who is using it. What we got till now: few hertz actual reads, high stat rate indicate of lots of people having the PATH and LD_LIBRARY_PATH set up.
- **lxplus5:** The current target date is 30 September 2014. These services have been replaced by SLC6-based services (lxplus6 and lxbatch6). Note that the deadline to report serious concerns for the termination of these SLC5-based services is 14 September 2014, please contact the CERN Service Desk (<http://cern.ch/service-portal> or service-desk@cern.ch).
- **Argus:** we suspect we've seen the same problem on the 28th (related to unresponsive CAs); so, we still think that fixes need to be properly released rather than manually installed from a github, and in addition it would be better if the problem was properly debugged/fixed.
- About lxplus5, Maite reports that there were only 15 tickets by individual users, not representing large experiment activities. Instructions were provided in the tickets.
- About Argus, Andrea Manzi suggests to update the existing ticket.

Tier 1 Feedback

- NDGF-T1 is testing FAX using native xrootd and nfs4-mount from dCache.

Tier 2 Feedback

Experiments Reports

ALICE

- low activity recently, new productions are being prepared
- CERN: job failure rates and efficiencies under investigation

ATLAS

- in addition to the reports of the specific task forces
 - ◆ DQ2 -> Rucio commissioning: FullChain test (T0-like workflow) ongoing, adding now automatic check of quotas. No direct implications for sites, but Rucio test and normal DQ2 production activity are producing a slightly higher load on the storage of the sites.
 - ◆ ProdSys1 -> ProdSys2 commissioning ongoing: many workflows implemented and tested. First "real production" workflows will start now. No implications for the sites
 - ◆ Activities:
 - ◇ Analysis as usual

- ◊ DC14 G4 simulation almost done
- ◊ DC14 mc14_13TeV reconstruction not started yet, pending validation, will run on multicore

- Alessandro reports that a few services are still using the AFS UI. The preferred date for decommissioning would be after the Quattor decommissioning, keeping it (at least in frozen state) until November-December.
- Alessandro reports that the planned retirement of lxplus5 is not an issue for ATLAS computing, and no issues reported yet by ATLAS offline.

CMS

- Processing overview:
 - ◆ Samples for CSA14 basically finished
 - ◆ Additional test workloads send to Tier-1 sites (mostly reading via xrootd from CERN EOS)
- SL5 lxplus5/lxbatch5
 - ◆ Computing coordinator in contact with Physics groups
 - ◊ Feedback will be given by mid of September
- AFS-UI at CERN
 - ◆ Many CMS SL5 based services and clients still use gLite clients from AFS
 - ◆ Efforts ongoing to move the services to SL6 (which has to happen anyhow)
 - ◆ Usage of Grid clients from CVMFS under investigation
- Reminders for sites
 - ◆ Upgrade CVMFS to 2.1.19
 - ◆ **Update xrootd fallback configuration**
 - ◊ Will start to open tickets - we keep on asking since months
 - ◊ Solid setup needed for AAA scale testing towards end of September
 - ◆ Add "Phedex Node Name" to site configuration
- Cristoph reports that keeping the AFS UI longer even in unattended mode would ease the transition for the remaining services.
- Alessandro asks to share the scale and expectations for the CMS xrootd federation (AAA) tests starting now - to be discussed offline.

LHCb

- LHCb activity dominated by simulation productions which, in waves, reach the full capacity.
- It has been x-checked that all references of DIRAC to the AFS UI have been removed. Nevertheless LHCb suggests to only rename the AFS UI volume in a first step in order to have a fallback in case things go wrong.
- Testing of voms3 client within LHCb environment will start soon.
- SHA2 certificate testing started. Many sites failing SAM tests and have been contacted by Maarten -> many thanks to him

Ongoing Task Forces and Working Groups

Network and Transfer Metrics WG

- Kick-off meeting [📅](#) will take place on Mon 8th of Sept at 3PM CEST
- Early version of the WLCG perfSONAR configuration interface will be deployed to production next week.
- Pythia Network Diagnosis Infrastructure (PuNDIT) project will be funded by NSF and starts at the beg. of September (lead by Shawn). The project will use perfSONAR-PS data to identify and localize

network problems using the Pythia algorithms. PuNDIT will collaborate with OSG and WLCG over its two year duration.

- Sites with incorrect versions of perfSONAR will receive tickets at the beg. of next week (9 sites in total)

WLCG perfSONAR service status report on 2014-09-04 10:03:22.949799 =====

```
perfSONAR instances monitored: 214
perfSONAR-PS versions deployed:
  3.2.2 : 6
  3.3.1 : 3
  3.3.2 : 173
  Unknown: 28
GOCDB registered total: 170
OIM registered total: 53
Unreachable instances (not monitored): 8
Incorrectly configured (failing >4 metrics): 28
```

- Andrea Sciaba asks if PuNDIT is also intended to also address the request to simplify the debugging of problematic links. Shawn answers that the project is about automated centralized analysis of monitoring data rather than debugging individual issues on links - but it can be used to raise alerts about specific problems.
- Maria comments that the dashboard linked in the WG documentation is password protected, Shawn answers that the protection can be removed. **Clarification from Shawn** *I misunderstood Maria's question. I thought she meant the Wiki page was protected. The OMD monitoring web interface requires a password as far as I understand. The MaDDash dashboard does not. I can explore options for the OMD (service checking component) to see if there is an alternative to requiring a password*
- Maria comments that it is difficult to use the dashboard to understand which links are commissioned. Shawn answers that PuNDIT is also intended to get attention on problematic points. **Addendum from Shawn** *We could institute something specific to tracking "links" (network paths). CMS has formally "commissioned" ~1600 links (40 sites squared) for their data transfer work. In effect MaDDash tracks source-destination behavior but not necessarily specific network routes. As noted PuNDIT should help in this area.*

Tracking Tools Evolution TF

FTS3 Deployment TF

- The FTS3 task force can be considered as successfully finished: FTS3 is now the production FTS service for all the LHC experiments.
- FTS3 overall: the plan is to have new releases every 3 - 4 months (unless urgent issues appear). We will organize dedicated demo/overview meetings in advance to demo and discuss the new features and functionality.
 - ◆ Just a reminder about communication:
 - ◇ For new feature requests, recommendations, general questions or open discussions
 - fts3-steering@cernNOSPAMPLEASE.ch (or pass by FTS3 dev office!)
 - ◇ For support and bug reporting, strictly technical issues
 - fts-support@cernNOSPAMPLEASE.ch
- Details of the past weeks:
 - ◆ FTS Dashboard:
 - ◇ New aggregation views have been included:
 - to monitor the active & queued transfers per endpoint/hosts
 - to monitor VO s activities (user files, production, T0 export) transfers
 - ◇ Ongoing tasks:
 - Fix the computing of transient state in FTS states monitoring view

- Implement the FTS3 VO/activity-shares monitoring

- Agreed to close the TF; applause.
- Maria reminds that there is also a GGUS support unit for FTS bug reports.

gLExec Deployment TF

- NTR

Machine/Job Features

- New contact for the condor implementation has been suggested by OSG and is currently being introduced to the topic -> welcome Marian Zvada.

Middleware Readiness WG

- We kindly asked last week T0 pre-prod machines to be equipped with the Package reporter in order to test scalability and to increase the packages database. Any news?
- The latest Cream-CE and Bdi update have been installed at LNL-T2 site (Legnaro). The configuration for the Job submission is ongoing (Andrea S is following this up)
- Design of the MW software for the verification activity is ongoing, a first prototype to be delivered before the next Meeting (1st October)
- Book your diaries!! Next meeting on October 1st at CERN with audioconf and vidyo. Provisional Agenda here [☞](#)
- About the deployment of the package reporter on the Tier-0, Maite answers that it can be followed up from this week.
- Andrea Manzi clarifies that the "MW software for the verification activity" uses the package reporter results to aggregate per software component, is used to tag good/bad versions, publishes the results in a dashboard.

Multicore Deployment

- **ATLAS:** 11 T1s and 35 T2s are currently running multicore for ATLAS with an average slot occupancy of 30k slots up to 57k. Among Tier2s production is still dominated by US dedicated sites but extra EU shared sites have been added in the past month. 1 French, 3 UK and 4 Italian Tier2s are among the sites added over the summer. Most of this sites have a dynamic or a mixture of dynamic and a small number of initial reserved slots. All sites are requested to enable dynamic batch system support to run multicore and single core jobs concurrently.
- **CMS:** Running multicore pilots regularly at all CMS T1s. Multicore pilots testing at several US T2s as well.
- **Passing parameters to the batch systems:** as reported in other occasions to apply certain scheduling techniques it is necessary for experiments to pass parameters to the batch systems ahead of being scheduled. To do this there have been many obstacles which need still to be cleared. One of these is the fact that there is not a standard blah script for sites deploying CREAM CE and those that exist are clearly not standardized since were written by sites trying to match what the experiments were sending but without any previous discussion. This has been discussed with ATLAS since there are other use cases for which this is needed for example requesting a specific amount of memory. It was decided the TF would take on board the standardization of the blah scripts (and other CEs scripts if needed) for the scheduling parameters. A discussion about distribution and deployment will also be needed.
- Alessandra Forti clarifies that the important point concerning blah script distribution is to have a

uniform choice of repository for all batch systems, it can be e.g. the WLCG repo or the CREAM product team repo.

- Agreed to assign the new task of standardizing the blah scripts to the multicore TF.

SHA-2 Migration TF

- introduction of the new VOMS servers
 - ◆ compliance of the WLCG infrastructure is being tested with the SAM **preprod** instances
 - ◇ ALICE [?](#) since July 23
 - ◇ LHCb [?](#) since Aug 22
 - ◇ ATLAS [?](#) and CMS [?](#) since Aug 28
 - ◆ EGI broadcast [?](#) sent Aug 25
 - ◆ deadline **Mon Sep 15**
 - ◇ switch SAM production configurations for **WLCG** on that day
 - ◇ gradually switch experiments' job and data systems as of that date
 - we need to get all relevant workflows verified **ASAP**
 - ◇ keep the new servers blocked for outside CERN a while longer?
 - ◆ 88 tickets opened on Sep 1 for EGI sites that still failed *due to the VOMS switch*
 - ◇ open tickets [?](#) (55 at the time of this meeting)
 - ◆ OSG sites are in very good shape
 - ◇ only WT2 (SLAC) SRM still failing *due to the VOMS switch*
 - ◆ EGI will run a campaign for Ops via the NGIs and ROCs
 - ◇ first try to ensure sites are configured OK
 - ◇ then switch SAM-Nagios instances
 - ◆ the old VOMS servers will continue running until the end of Nov
 - ◇ their host certs expire by then
- Alessandro asks how the experiments can test the new VOMS servers in their services; Maarten answers that the config needs to be changed on the machines contacting VOMS as done for SAM preproduction.
- Nicolo' reminds that many services use the AFS UI to contact VOMS, so the config will also need to be changed there. To be done only when well tested.

WMS Decommissioning TF

- Condor-based SAM probes
 - ◆ Deployment to production planned on Wed 1st of October 2014
- Maria asks if the SAM BDII can also be decommissioned. Maarten reminds that the new CE probes (CREAM and Condor) still need the BDII and it's better to have a dedicated one, so the answer is no..

IPv6 Validation and Deployment TF

- Nothing to report.

Squid Monitoring and HTTP Proxy Discovery TFs

- Automated MRTG monitor [?](#) needs a few tweaks before we can start registration campaign, in particular
 1. Use site names when reading from OIM instead of service names
 2. Distinguish between multiple squid services registered at a site
 3. Use different service type from OIM after OIM is changed to not distinguish between Frontier & CVMFS squids

- Made a draft of a WLCGSquidRegistration instruction page. Will definitely need more work.
- Agreed in a meeting to incorporate Costin Grigoras' new MonALISA-based monitor into the wlcg-squid-monitor as an additional interface and to use his rpm-based host collector as an extra optional method to collect extra monitoring data. So far we still expect every site to also make their squid monitoring accessible via SNMP.
- Dave comments that the issue with squid registration in OIM does not affect GOCDB (where the proper registration was requested in the first place)
- Alessandro asks if the sites need to change anything for the new optional monitoring. Dave answers no: the new monitoring can also use SNMP data, but it will also use additional info if published by the site.

Action list

1. ONGOING on the WLCG middleware officer and the experiment representative: for the experiments to report their usage of the AFS UI and for the middleware officer to take the steps needed to enable the CVMFS UI distribution as a viable replacement for the AFS UI.
2. ONGOING on Andrea S.: to understand with EGI if it is possible to bypass the validation of VO card changes in the case of the LHC VOs. Agreed by Peter Solagna.
3. ONGOING on the WLCG monitoring team: status of the CondorG probes for SAM to be able to decommission SAM WMS. Status: see these minutes.
4. CLOSED on the middleware officer: report about progress in CVMFS 2.1.19 client deployment. Status: only 2 T3s missing.
5. ONGOING on the WLCG monitoring team: evaluate whether SAM works fine with HTCondor CE. Report about showstoppers. Status: the SAM team made a proposal on the steps to taken to enable SAM. ATLAS is following up to make sure that the new CEs are correctly visible in AGIS, while for the CMS VO feed they will be taken directly from OIM. The plan is at first to test HTCondor-CEs in preproduction and later switch to production. It is not foreseen to monitor at the same time GT5 and HTCondor endpoints on the same host.
 - ♦ Maarten comments that from private communication with Marian, SAM seems fine after the arrangement with OSG to hack the information needed for discovery into OIM.
6. NEW on Alessandro DG: find out from OSG about plans for publication of HTCondor CE in information system, and report findings to WLCG Ops. To be followed up with Michael Ernst and Brian Bockelman.

AOB

- Simone asks to clarify the questions to the experiment representatives posted in the agenda of the next GDB. In particular about "metadata protocol"
- Michel clarifies that "metadata protocol" means "the protocol (e.g. SRM, xrootd, WebDAV) used to manipulate metadata on the storage". He will reformulate the questions in the agenda.
- Alessandro reminds that ATLAS presented the plans for the usage of the protocols at many meetings, the final choice of protocol will also depend on the measured performances.
- Michel explains that the point of the questions is also to understand if there have been new developments based on the tests performed since the plans were announced a few months ago.

-- MariaALANDESPRADILLO - 20 Jun 2014

This topic: LCG > WLCGOpsMinutes140904

Topic revision: r28 - 2014-09-05 - ShawnMcKee



Copyright &© 2008-2024 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

WLCGOpsMinutes140904 < LCG < TWiki

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback