

Table of Contents

WLCG Operations Coordination Minutes - December 5, 2013.....	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
Experiments operations review and Plans.....	4
ALICE.....	4
ATLAS.....	4
CMS.....	5
LHCb.....	5
Discussion on Experiment Plans.....	5
Ongoing Task Forces Review.....	5
WMS decommissioning.....	5
gLExec.....	6
FTS3.....	6
Tracking tools evolution.....	6
perfSONAR.....	6
IPv6.....	6
Middleware readiness.....	7
Machine/Job Features.....	7
SHA-2.....	7
Action list.....	8

WLCG Operations Coordination Minutes - December 5, 2013

Agenda

- <https://indico.cern.ch/conferenceDisplay.py?confId=280059>

Attendance

- Local: Andrea Sciabà (chair), Simone Campana (secretary), Maarten Litmaath, Vincent Brillault, Oliver Keeble, Alberto Aimar, Ivan Glushkov, Michail Salichos, Nicolò, Felix Lee, Markus Schulz, Maite Barroso, Maria Dimou, Alessandro Di Girolamo
- Remote: Joel Closier, Shawn Mc Kee, Antonio Perez Calero Yzquierdo, Renaud Vernet, Thomas Hartmann, Christoph Wissing, Valery Mitsyn, Frederique Chollet, Alessandro Cavalli, Alessandra Doria, Gareth Smith, Borut Kersevan, Alastair Dewhurst, Andrea Valassi, Isidro Gonzalez Caballero, Alessandra Forti

News

- Following the discussions at the last planning meeting on multicore resource deployment, we would like to propose a dedicated task force
 - ♦ the draft of the mandate is available
 - ♦ exceptionally we should use this meeting to discuss the proposal and decide whether the task force should be created
 - ♦ **Discussion** Andrea presented the draft mandate of the task force. There were some concerns (Maite, Joel) expressed about the overlap of the TF with the Machine/Job Features TF and the Cloud WG. Andrea/Simone pointed out that the overlap exists in some corners of the various initiatives but for a small fraction, so it does still make sense to keep the initiatives separate, while trying to minimize the overlap. Alessandra/Antonio commented that the fact that several members participate in many of those initiatives facilitates the process of reducing the overlap. Maarten suggested that the situation should be reviewed after some time and initiatives could be merged if we feel there is larger overlap than expected. The **Multi Core Deployment Task Force has been approved**. Alessandra Forti and Antonio Perez-Calero will lead it.
- Experiment plans and needs during the Christmas period are discussed below.

Middleware news and baseline versions

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

- **Discussion** Maarten commented that the update of the baseline version of the BDII is particularly important for the SAM BDII nodes at CERN, because the new version will finally rid those nodes of the FCR mechanism that should never have been running there in the first place! Will be followed up.

Tier-1 Grid services

Storage deployment

Site	Status	Recent changes	Planned changes
CERN			

Site	Status	Recent changes	Planned changes
	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 0.3.4 / xrootd 3.3.4 / BeStMan2-2.2.2) CMS (EOS 0.3.2 / xrootd 3.3.4 / BeStMan2-2.2.2) LHCb (EOS 0.3.3 / xrootd 3.3.4 / BeStMan2-2.2.2)		
ASGC	CASTOR 2.1.13-9 CASTOR SRM 2.11-2 DPM 1.8.7-3 xrootd 3.3.4-1	DPM 1.8.6-1 --> 1.8.7-3, DPM xrootd 3.2.7-1 --> 3.3.4-1	
BNL	dCache 2.2.10 (Chimera, Postgres 9 w/ hot backup) http (aria2c) and xrootd/Scalla on each pool		dCache upgrade to v2.6 on Dec 17 for SHA-2 compatibility
CNAF	StoRM 1.11.2 emi3 (ATLAS, CMS, LHCb)	none	none
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 Oracle Lustre 1.8.6 EOS 0.3.2-4/xrootd 3.3.3-1.slc5 with Bestman 2.2.2.0.10		Will upgrade xrootd/EOS after next EOS release; a dCache 2.2 pool is up and we are starting the transition process
IN2P3	dCache 2.6.15-1 (Chimera) on SL6 core servers and pool nodes Postgres 9.2 xrootd 3.0.4		
KISTI	xrootd v3.2.6 on SL5 for disk pools xrootd 20100510-1509_dbg on SL6 for tape pool dpm 1.8.7-3	DPM 1.8.6-1 --> 1.8.7-3	
KIT	dCache <ul style="list-style-type: none"> atlassrm-fzk.gridka.de: 2.6.5-1 cmssrm-fzk.gridka.de: 2.6.5-1 lhcb-srm-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 		We want to update all dCache setups to at least 2.6.15 this year. LHCb was already upgraded
NDGF			

Site	Status	Recent changes	Planned changes
	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 2.2.17-1 xrootd door to VO servers (3.3.1-1)		
RAL	CASTOR 2.1.13-9 2.1.13-9 (tape servers) SRM 2.11-1		
TRIUMF	dCache 2.2.18		

- **Discussion** The table probably is not accurate for the CERN EOS versions. Instances of all experiments should now be at the latest version. Will be verified and fixed if needed.

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		
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	1.8.3.1-1 for T1 and US T2s	SL6, gLite	ORACLE 11gR2	ATLAS	
CERN	1.8.7-3	SLC6, EPEL	Oracle 11	ATLAS, LHCb, OPS, ATLAS Xroot federations	

Other site news

Status of Tier-1 WN deployment on OPN now tracked in this survey:

- WNs on the OPN

NOTE: this is NOT a request for deployment, it is a survey of the current status to facilitate experiment operations planning.

Experiments operations review and Plans

ALICE

- plans for the end-of-year break
 - ◆ MC production at all sites
 - ◆ we do not expect to run RAW reconstruction
 - ◆ the user/organized analysis will naturally diminish in intensity
 - ◆ the usual 'best effort' support from the sites, which worked so well in the past years, will be appreciated!
- CERN
 - ◆ SLC6 vs. SLC5 job failure rates and CPU/wall-time efficiencies
 - ◊ 4 VOBOXes submit to CERN resources since a week
 - CERN-L: SLC5 + Torrent + Meyrin
 - CERN-CREAM: SLC5 + CVMFS + Meyrin
 - CERN-SHA2: SLC6 + CVMFS + Meyrin
 - CERN-CVMFS: SLC6 + CVMFS + Wigner
 - ◊ longer comparison time needed to average out effects due to different job types
 - to be continued
- CVMFS
 - ◆ 61 sites using it in production
 - ◆ 11 in various stages of preparation
 - ◆ sites please ensure the WN have version 2.1.15 (or higher)
- SAM
 - ◆ ALICE sites that care about the SAM Availability and Reliability reports should ensure they look OK in the ALICE SAM tests, which will be used instead of the Ops tests as of Jan 2014:
 - ◊ ALICE SUM Dashboard [↗](#)
 - ◊ ALICE MyWLCG summary [↗](#)
 - ◊ ALICE Nagios page [↗](#)
- **Discussion** Andrea asked if there is coordination between Alice and ATLAS concerning the CPU/WCT tests at CERN. Maarten: the tests from the two experiments look at complementary aspects but there is cross communication of the results and issues found.

ATLAS

- ATLAS holiday break plans
 - ◆ MC production: we have to produce 130M events starting from the next days . This corresponds to approx 10 days of ATLAS Grid production resources utilization. More tasks are being defined in these days
 - ◆ Reprocessing: a reprocessing campaigns is foreseen to start in the next week. 2.2PB of inputs on tape, small output (2%). This corresponds to approx 30 days for 20% of the T1s. Pre-stage of the data is automatically handled by PanDA, but we will produce the list of data for each Tier1 and we will advertise to them.
 - ◆ Group prod: NTUP_COMMON campaign now freezing the code. by 16th of December validation of the slice tests. If everything goes well it will start before Christmas. It corresponds to 35% of all the resources for approx 5 weeks.
 - ◆ analysis as usual
 - ◆ more news in 2 weeks from now.
- as for CMS: Best-effort operations during holiday break as every year
 - ◆ Appreciate all support from the sites we can get
 - ◆ Will still send tickets as usual.

CMS

- CMS holiday break plans:
 - ◆ Production and digitization-reconstruction of Run 2 preparation MC samples
 - ◆ Digitization-reconstruction of 7 TeV MC for 2011 data legacy re-reconstruction pass
- Best-effort operations during holiday break as every year
 - ◆ Appreciate all support from the sites we can get, but don't expect normal levels of support, especially for T2 sites
 - ◆ Will still send tickets though
- Multicore deployment task force
 - ◆ Important is the coordination with the machine Machine / Job Features Task Force and the Cloud group
 - ◆ CMS representative will be Antonio Perez-Calero Yzquierdo from PIC

LHCb

- Fall '13 incremental stripping campaign is finished
 - ◆ Main operation finished within 6 weeks which was the minimum scheduled time for it. Thanks to the excellent performance of all T1 sites.
 - ◆ Remaining few missing files (O(10s)) are looked after
- Plans during Xmas shutdown
 - ◆ Usage of distributed grid resources for mainly monte carlo productions. Surveillance by the operations team on a best effort basis.
- Downtime of Dirac services scheduled for Monday Dec 9th morning, b/c of a physical move of hardware needed in the CC
 - ◆ During this migration several Dirac agents will move to VMs and some databases will be migrated to the DB on demand service
- Grid aware directories under EOS have been moved to a new directory structure to allow also migration of "castor user files"
 - ◆ grid aware user files will be migrated from CASTOR to EOS during downtime on 9 Dec
 - ◆ migration for the remaining castor user files will be done next year
- Reminder for sl6 migration, LHCb will stop building slc5 binaries as of January 2014 for its application stack
- New dashboard page
<http://dashb-lhcb-ssb.cern.ch/dashboard/request.py/siteview#currentView=CVMFS> available which visualizes the status of the CVMFS probe run by nagios

Discussion on Experiment Plans

- **All experiments will run activities over christmas at non negligible scale. They do not require special effort from sites or WLCG in general, while best effort support is highly appreciated.**
- CERN was asked about plans for SLC6 deployment on WNs and plans for upgrading CVMFS clients to the latest version. Maite: the initial target was 100% of the WNs by the end of the year, but this had to be revisited. The new target is 100% by the end of January 2014. For CVMFS clients, this is bound to the SLC6 upgrade: SLC6 nodes do come with the new client. There is no point investing effort in upgrading SLC5 since the complete migration to to SLC6 will finish in a bit more than a month.

Ongoing Task Forces Review

WMS decommissioning

- usage of the CMS WMS at CERN seems to have gone down since CMS users were informed that support of the gLite WMS is ramping down and they should use CRAB's

`scheduler=remoteglides` option instead

- ◆ the CRAB-2 client also no longer uses a centrally distributed list of WMS hosts

- **Discussion** Maarten explained that the CMS usage pattern of WMS nodes at CERN goes in spikes. Generally it went down, but every once in a while there are bursts of activity. In the next weeks (probably not this year) we will try to get the remaining usage ramped down further and then decommission those nodes at a time that is not too inconvenient for CMS.

gLExec

- 61 tickets closed and verified, 33 still open
 - ◆ some sites still waiting to finish their SL6 migration first
 - ◆ some difficult cases being debugged
- 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 at IN2P3-SUBATECH has led to even stranger errors
 - ◇ ALICE Nagios page [↗](#)
 - ◆ to be followed up
- Deployment tracking page
- **Discussion** Andrea asked if it would be possible to remove the dependency on the perl module `Time/HiRes.pm` from the probe. Maarten clarified that this would basically mean rewriting the probe, while the current probe is quite well written and works fine at most sites.

FTS3

- specific FTS3 link performance tests with autoconf vs fixed conf are in progress.

Tracking tools evolution

- GGUS availability notes for Year End are here: *For the Year End period: GGUS is monitored by a monitoring system which is connected to the on-call service. In case of total GGUS unavailability the on-call engineer (OCE) at KIT will be informed and will take appropriate action. Apart from that WLCG should submit an alarm ticket which triggers a phone call to the OCE.*
- Very useful meeting at KIT between Christoph, Helmut and Guenter clarified a number of important development items for the elimination of the savannah-GGUS bridge. Minutes in Savannah:131565#comment34 [↗](#)

perfSONAR

- See slides
- **Discussion** Andrea asked what is the plan to find manpower for development/maintenance of the perfSONAR dashboard code. Simone asked if it was possible to try with BNL since the previous developer was from there. Shawn explained that most likely BNL would not be able to participate in the development effort because of cuts in the funding. A possible solution is being looked into between OSG and ESNNet.

IPv6

- See slides
- Andrea/Simone: the request of CMS to have IPV6 supported also on SLC5 at CERN has been already discussed with Edoardo Martelli. The request comes from other groups as well and Edoardo agrees on providing a solution. He will work on it next week.

Middleware readiness

- first meeting on Dec 12 2013 at 16hrs CET with audioconf possibilities. Agenda <https://indico.cern.ch/conferenceDisplay.py?confId=285681>

Machine/Job Features

- MachineJob Features meeting with Igor Sfiligoi on his CHEP presentation working on minimizing draining waste (MDW) cpu time for multi-core pilots
 - ◆ The proposal includes a bi-directional communication between pilots and resource providers whereas MachineJobFeatures only proposes uni-directional communication resource->pilot
 - ◆ MDW and MJF are very similar for the communication resource->pilot and will try to include the missing bits of the other proposal into their own
 - ◆ The communication pilot->resource will be investigated by MJF, it is a possibility to provide more detailed information to the resource on which pilots to stop in case this is needed
- With a common approach we will both profit from having all batch system types covered for both use cases.

SHA-2

- sites are steadily upgrading remaining affected services to versions supporting SHA-2
 - ◆ SHA-2 migration update in Nov 28 EGI OMB meeting mentioned 5 dCache and 7 StoRM instances still to be done
 - ◆ OSG T1 sites
 - ◇ BNL plan for Dec 17
 - ◇ FNAL hopefully OK by the end of Dec
- client issue reported by OSG interoperability liaison Anthony Tiradani:
 - ◆ dCache SRM client needs a newer version (2.6.12) to be able to handle SHA-2 **host** certificates!
 - ◇ released on Dec 2 as part of EMI-3 Update 11
 - ◇ not really used by the experiments?
 - Christoph: some CMS sites may have configured that client for their local staging operations; some might need the upgraded client still on EMI-2 then
 - Maarten: will follow up
- the experiments have tested a lot and look ready
- timelines
 - ◆ by mid January the WLCG infrastructure is expected to be essentially ready
 - ◇ we may be able to ignore any remaining stragglers by the end of Jan
 - ◆ it is unlikely for SHA-2 certs to appear still this year
 - ◇ the OSG CA foresees starting mid Jan
 - ◇ the CERN CA will switch when WLCG is ready
- VOMRS
 - ◆ a VOMS-Admin test setup was successfully loaded with the VOMRS data of ALICE two weeks ago
 - ◆ the setup needed to be redone after its configuration was cleaned up
 - ◆ preparation of the upgrade of the production VOMS services and fixing the subsequent fallout took a lot of time
 - ◆ the test setup is still to become available soon
 - ◇ VOMS-Admin instability being investigated...
- **Discussion** Christoph pointed out that the baseline version of grid clients is still EMI-2 while SHA-2 support for dCache clients comes with EMI-3. Maarten asked if CMS still needs the dCache clients. Christoph explained that while they are not broadly used, there are Data Management agents deployed at the sites where, depending on local customizations, the dCache clients may be used. After

reminding that there currently is a problem with EMI-2 WN installations and updates on SL6 (dependency error), Maarten agreed that there is a use case then for making a SHA-2 compliant dCache client available also in EMI-2 and he will discuss this further with the developers.

Action list

1. Tracking tools TF members who own savannah projects to list them and submit them to the savannah and jira developers if they wish to migrate them to jira. AndreaV and MariaD to report on their experience from the migration of their own savannah trackers. Further discussion expected for the next meeting, after a dedicated meeting about the migration of the GGUS Savannah tracker to JIRA). Maria clarifies that 83 trackers need a decision, and trackers that will not be migrated will be gone for good. Maarten suspects the migration cannot be finished this year, but will need to stretch a few months into next year. **Maria thinks we can close this action as decided in the Tracking Tools' Evolution TF meeting of 2013/10/08. See Minutes [HERE](#). closed**
2. Investigate how to separate Disk and Tape services in GOCDB
 - ◆ proposal submitted via GGUS:93966 [↗](#)
 - ◆ **in progress** - ticket updated
3. Agree with IT-PES on a clear timeline to migrate OPS and the LHC VOs from VOMRS to VOMS-Admin
 - ◆ **in progress**
4. Experiments interested in using WebDAV should contact ATLAS to organise a common discussion
 - ◆ **closed**
5. Collect feedback from VOs about need for grid-cert-info and setting EMI-UI 2.0.3 as baseline.
 - ◆ **new**

-- SimoneCampana - 27 Nov 2013

This topic: LCG > WLCGOpsMinutes131205

Topic revision: r31 - 2018-02-28 - MaartenLitmaath



Copyright &© 2008-2024 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