

Table of Contents

WLCG Operations Coordination Minutes - December 19, 2013.....	1
Agenda.....	1
Attendance.....	1
News.....	1
Middleware news and baseline versions.....	1
Tier-1 Grid services.....	1
Storage deployment.....	2
FTS deployment.....	3
LFC deployment.....	3
Other site news.....	4
Experiments operations review and Plans.....	4
ALICE.....	4
ATLAS.....	4
CMS.....	6
LHCb.....	6
Ongoing Task Forces and Working Groups.....	6
SHA-2 Migration TF.....	6
Tracking tools evolution TF.....	7
gLExec deployment TF.....	7
Middleware readiness WG.....	7
WMS decommissioning TE.....	7
FTS3 Deployment TE.....	8
Action list.....	8

WLCG Operations Coordination Minutes - December 19, 2013

Agenda

- <https://indico.cern.ch/conferenceDisplay.py?confId=282475>

Attendance

- Local: Andrea Sciabà, Felix Lee, Maarten Litmaath, Stefan Roiser, Jerome Belleman, Alessandro Di Girolamo, Pablo Saiz, Oliver Keeble, Nicolò Magini, Maite Barroso Lopez
- Remote: Shawn Mc Kee, Burt Holzman, Giovanni Zizzi, Oliver Gutsche, Renaud Vernet, Massimo Sgaravatto, Frederique Chollet, Di Qing, Thomas Hartmann, Antonio Maria Perez Calero Yzquierdo, Rob Quick, Michel Jouvin, Alexey Sedov, Gareth Smith, Joao Pina

News

There was no special news this time.

Middleware news and baseline versions

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

In response to the OpenSSL issue also the minimum GridSite versions are listed explicitly now. Maarten explained that the OpenSSL matter is not fully understood at this time:

- there has **not** been a big impact on the infrastructure so far
- in direct job submission tests with CERN CREAM and UI instances the delegated proxies ended up with 1024-bit keys, despite that nothing was updated
 - ◆ but 512-bit keys can still be reproduced at DESY-ZN
- the SAM WMS have **not** been updated with the new `gridsite` version, hence continue generating 512-bit proxies, yet nobody reported problems due to that
 - ◆ the new `gridsite` has been tested and the update can be done at short notice, if needed
 - ◆ otherwise it will be done in Jan
- sites are advised to keep SL6 services on SL6.4 for the next 2 weeks, unless an urgent security update requires SL6.5

During the preceding operations meeting Rob reported that OSG have done an emergency release on Tue to fix the affected Globus components

- there were complications due to the use of a private interface to OpenSSL; the code now uses a public interface instead

The deadline for sites to upgrade CVMFS from v2.0.x to v2.1.15 or higher is March 1, to allow the Stratum-0 service to be upgraded to v2.1 at last, which cannot be done while v2.0 clients still need to be supported. Some of the experiments foresee adding CVMFS tests to their critical SAM profiles at some point.

Tier-1 Grid services

Storage deployment

Site	Status	Recent changes	Planned changes
CERN	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	none	none
BNL	dCache 2.6.18 (Chimera, Postgres 9 w/ hot backup) http (aria2c) and xrootd/Scalla on each pool	dCache upgrade to v2.6 for SHA-2 compliance	
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 EOS 0.3.2-4/xrootd 3.3.3-1.slc5 with Bestman 2.2.2.0.10	Lustre decommissioned in favor of EOS	Will upgrade xrootd/EOS after next EOS release (if FUSE bugs are fixed); a dCache-disk pool (Chimera + dCache 2.2) is up, links are being commissioned and all data to be migrated is being pinned to the existing dCache 1.9 instance
IN2P3	dCache 2.6.18-1 (Chimera) on SL6 core servers and pool nodes Postgres 9.2 xrootd 3.0.4	dCache 2.6.15-1 --> 2.6.18-1	xrootd 3.3.4
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"> • atlasrm-fzk.gridka.de: 2.6.17-1 • cmssrm-kit.gridka.de: 2.6.17-1 • lhcbstrm-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 		
NDGF	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 (3.3.4)		today upgrade DBs to PostgreSQL9.2 & upgrade of dCache to 2.2.21
RAL	CASTOR 2.1.13-9 2.1.13-9 (tape servers) SRM 2.11-1		CASTOR 2.1.14 in testing
TRIUMF	dCache 2.2.18		
JINR-T1	dCache <ul style="list-style-type: none"> • srm-cms.jinr-t1.ru: 2.6.19 • srm-cms-mss.jinr-t1.ru: 2.2.23 with Enstore xrootd federation host for CMS: 3.3.3		

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	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		
JINR-T1	2.2.8 - transfer-fts-3.7.12-1		We plan to install additional fts3 on separate host

LFC deployment

Site	Version	OS, distribution	Backend	WLCG VOs	Upgrade plans
BNL		SL6, gLite		ATLAS	None

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

Other site news

Experiments operations review and Plans

ALICE

- plans for the end-of-year break (reminder)
 - ◆ 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 different sets of CERN resources since 3 weeks
 - ◇ Wigner job failure rate was 55% compared to 18% for SLC6 jobs in Meyrin and 30% for SLC5 jobs
 - ◇ the average efficiency of SLC6 jobs was 20% lower than the average for SLC5 jobs
 - ◇ similar comparisons for various classes of ATLAS and CMS jobs suggest differences ranging from 0 to 20% depending on the type of job
 - ◇ a queue targeting only physical SLC6 nodes would help to understand if the differences are due to SLC6 or due to the VM infrastructure
 - ◇ to be continued...
- RRC-KI-T1
 - ◆ commissioning activities ongoing since late Nov - thanks!
 - ◇ EOS, VOBOX, CEs
- CVMFS
 - ◆ 64 sites using it in production
 - ◆ 8 in various stages of preparation
 - ◆ sites please ensure the WN have version 2.1.15 (or higher)
- SAM (reminder)
 - ◆ 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 [↗](#)

Alessandro noted it would be easier to investigate job efficiency issues and failure rates at CERN if there were queues with uniform resources. Maite replied that this can be discussed further in the ad-hoc working group dedicated to investigating these matters. Stefan said that also LHCb may join in, though there have been no concerns about the LHCb job performance on SLC6 so far.

ATLAS

- ATLAS xmas break plans
 - ◆ MC production single job: produce 100M events, which corresponds to approx 8/10 days of ATLAS Grid production resources utilization as of today.

- ◆ MC production MultiCore: produce 150M events, tasks are being tested now. If everything goes well as expected, the message to sites is:
 - ◇ Configure MCore queues fully dynamically if experienced with it,
 - ◇ Static allocation otherwise: 50% of production resources for T1 and big T2
 - ◇ if for some reason the MultiCore configuration is not production ready at the site, do not increase the share now (before/during XMas) if you think the system stability could be endangered. Please communicate with ATLAS on when you think you can do these changes.
- ◆ Reprocessing: a reprocessing campaigns has started. The total of the inputs is 2.2PB 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. During XMas Break only a part of it is foreseen to be done (approx 500TB of inputs)
- ◆ Group prod: NTUP_COMMON v2 campaign is now starting. It corresponds to 35% of all the resources for approx 5 weeks.
- ◆ analysis as usual
- ◆ check more details on Tuesday 17 December ADC Weekly agenda [↗](#)
- Issues: Openssl
 - ◆ for sites and central services, please read [https://twiki.cern.ch/twiki/bin/view/LCG/WLCGDailyMeetingsWeek131216#Monday "openssl issue"](https://twiki.cern.ch/twiki/bin/view/LCG/WLCGDailyMeetingsWeek131216#Monday%20%22openssl%20issue%22)
- information for sites
 - ◆ Rucio Renaming: deadline 1st of February. sites not migrated (or not started/agreed) will be excluded from ATLAS DDM. ADC Weekly 17 December - Rucio Renaming: Deadline [↗](#)
 - ◇ Exceptions can be discussed for sites with migration in progress with clear plans agreed beforehand with DDM/Rucio teams
 - ◇ What we expect from the not yet renamed sites:
 - DPM and dCache sites must provide a WebDAV access before this date to allow remote renaming. If they cannot/do not want, they have to contact atlas-dq2-ops@cern.ch and they will have to take care of the renaming themselves.
 - StoRM sites: we notice that the performance of the current implementation of WebDAV is not good enough. StoRM developers are working on an improved version, but it might be tight to have it deployed on all StoRM sites by February 1st. The sites have to be ready to upgrade their storage as soon as possible (e.g. the beginning of January) if StoRM release will be ready as expected.
 - ◆ MultiCore allocations for production (as described above under XMas activity) :
 - ◇ Configure MultiCore queues fully dynamically if experienced with it,
 - ◇ Static allocation otherwise: 50% of production resources for T1 and big T2
 - ◇ if for some reason the MultiCore configuration is not production ready at the site, do not increase the share now (before/during XMas) if you think the system stability could be endangered. Please communicate with ATLAS on when you think you can do these changes.

Giovanni said the next version of StoRM is expected in a few days and that after successful testing first the INFN-T1 would be upgraded, hopefully by mid Jan, before the release would be proposed to T2 sites. Alessandro noted the timeline may have an impact on the Rucio plans.

Antonio asked what is requested from the sites w.r.t. the multi-core deployment: partition the WN, have separate queues? Alessandro replied that the resources need not be partitioned, but that separate queues are needed indeed. Sites can further discuss the details via their cloud support team etc.

CMS

- Reminder 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
- Reminder: 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

LHCb

- Plans for the Xmas break
 - ◆ Started a ProtonIon/IonProton reprocessing.
 - ◇ RAW data staged in at CERN & GRIDKA. CERN all data is staged, GRIDKA progressing well
 - ◇ bulk of the work shall be finished in ~ 10 days
 - ◆ Otherwise Monte Carlo productions at all Sites / Tier levels
- GGUS statistics: 17 tickets opened in last 2 weeks, mainly problems with pilots aborted (9)
- A fix for resolving of the xroot address by SRM correctly on DPM sites has been tested successfully at CBPF.
- LHCb hit by problem caused by latest openssl version on redhat linux versions
 - ◆ Found also for FTS3 transfers. Mitigated by CERN by rolling back to a previous version.
 - ◆ Info for sites at: <https://operations-portal.egi.eu/broadcast/archive/id/1066>

Oliver said the DPM fix will be released in Jan.

Ongoing Task Forces and Working Groups

SHA-2 Migration TF

- sites are steadily upgrading remaining affected services to versions supporting SHA-2
 - ◆ Operations update in Dec 19 EGI OMB meeting mentioned 6 sites with non-compliant services remaining
 - ◆ OSG T1 sites
 - ◇ BNL ready since Dec 17
 - ◇ FNAL hopefully OK by the end of Dec
 - cmssrmdisk.fnal.gov seems OK
 - cmssrm.fnal.gov not yet ready
- EOS SRM instances **not** yet ready!
 - ◆ updated version tested **OK** on eospps.cern.ch and standby nodes for the experiments
 - ◇ can be switched quickly if needed
 - ◆ updates of the production instances early Jan
- newer dCache SRM client v2.2.22 able to handle SHA-2 host certificates released on Dec 16
 - ◆ now based on Java 6 instead of 7
 - ◆ EMI-2 Update 21
 - ◆ EMI-3 Update 12
- 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
 - ◆ VOMS-Admin test setup should become available for testing by VO managers early Jan
 - ◆ VOMS-Admin instability being looked into (GGUS:99327) [↗](#)
 - ◇ thanks to the VOMS developers for their prompt efforts!

Burt reported the following after the meeting:

We're not going to make the end of December and even January is optimistic -- I think we are realistically looking at mid to late February for all of the Tier 1 storage elements to be migrated. The reason for the delay is mostly not technical -- but has to do with the fact that the CMS LPC storage is also on dCache, and we want to give our analysis users sufficient time to migrate their data to EOS. The goal at this point is to give users a hard deadline of January 31 before we are able to decommission the old SE.

Maarten notes that this matter looks mostly internal to FNAL and CMS.

Tracking tools evolution TF

- GGUS: Reminder: For the Year End period: GGUS is monitored by a system 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. If GGUS is available but there is a problem with the workflow, e.g. ALARM to CERN doesn't generate email notification to the operators, then WLCG should submit an ALARM ticket, notifying Site DE-KIT, which triggers a phone call to the OCE. If the web portal is unavailable, contact details for KIT are recorded in the GOCDB.

gLExec deployment TF

- 64 tickets closed and verified, 31 still open
 - ◆ some sites still waiting to finish their SL6 migration first
 - ◆ progress for 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 now looks sufficient to cure the problem
 - ◆ a proper fix is still to be decided
- Deployment tracking page

Maarten confirmed there are sites that have finished migrating to SL6, but not yet managed to get their gLExec infrastructure working. He recalled that CMS intend to make gLExec tests critical early next year, so at least the CMS sites should pay attention there as of Jan.

Middleware readiness WG

- first meeting [↗](#) happened on Dec 12
 - ◆ the discussion was mostly on **repositories** and **processes**
 - ◇ how to involve experiments and sites should be discussed in the next meeting, which is planned for Feb 6 at 16:00 CET
 - ◆ please consult the minutes [↗](#) for a detailed summary

WMS decommissioning TF

- usage of the CMS WMS at CERN has remained lower since CMS users were informed that support of the gLite WMS is ramping down and they should use CRAB's `scheduler=remoteglitein` option instead
 - ◆ the CRAB-2 client also no longer uses a centrally distributed list of WMS hosts

- to be continued after the break

Andrea reported he informed the Geant4 VO of the decommissioning plans and showed them how to find other WMS instances supporting the VO. Maarten added that migration strategies have to be found for each of the VOs still relying on the CERN instances today.

FTS3 Deployment TF

- FTS3 servers affected by issue with new openssl version, for now rolled back to SLC6.4, permanent fix when new gridsite version is released in EPEL
- FTS3 performance comparison (fixed conf vs autoconf) tests ongoing - some bugs reported to developers; collecting preliminary results on FTS3 twiki

Action list

1. Investigate how to separate Disk and Tape services in GOCDB
 - ◆ proposal submitted via GGUS:93966 [↗](#)
 - ◆ **in progress** - ticket updated
2. Agree with IT-PES on a clear timeline to migrate OPS and the LHC VOs from VOMRS to VOMS-Admin
 - ◆ **in progress**
3. Collect feedback from VOs about need for grid-cert-info and setting EMI-UI 2.0.3 as baseline.
 - ◆ **closed**

It was agreed that the last action item can just be closed.

-- SimoneCampana - 12 Dec 2013

This topic: LCG > WLCGOpsMinutes131219

Topic revision: r27 - 2013-12-20 - MaartenLitmaath



Copyright &© 2008-2021 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