

Table of Contents

WLCG Operations Coordination Minutes - 18 July 2013.....	1
Agenda.....	1
Attendance.....	1
News.....	1
Middleware news and baseline versions.....	2
Tier-1 Grid services.....	2
Storage deployment.....	2
FTS deployment.....	3
LFC deployment.....	4
Other site news.....	4
Data management provider news.....	4
Experiment operations review and plans.....	4
ALICE.....	4
ATLAS.....	5
CMS.....	5
LHCb.....	6
News from EGI Operations.....	6
Task Force reports.....	7
SL6.....	7
SHA-2 migration.....	7
FTS-3.....	7
gLExec.....	8
perfSONAR.....	8
AOB.....	8
Action list.....	8

WLCG Operations Coordination Minutes - 18 July 2013

Agenda

<https://indico.cern.ch/conferenceDisplay.py?confId=260737>

Attendance

- Local: Maria Girone (chair), Andrea Sciabà (secretary), Maarten Litmaath, Oliver Keeble, Alessandro Di Girolamo, Helge Meinhard, Jan Iven, Domenico Giordano, Ian Fisk, Ikuo Ueda, Nicolò Magini
- Remote: Joel Closier, Vanessa Hamar, Ilya Lyalin, Oliver Gutsche, Alexander Verkooijen, Massimo Sgaravatto, Felix Lee, Alessandra Doria, Jeremy Coles, Di Qing, Peter Solagna, Sang Un Ahn, Gareth Smith, Stephen Burke, Alessandra Forti, Thomas Hartmann, Josep Flix

News

Maria proposes to have the next meeting on August 29 and September 19, due to the several absences in August and September 5 being a holiday at CERN. The proposal is approved.

Maarten:

- Middleware production readiness verification task force
 - ◆ See the presentation [by Markus Schulz](#) in the July 16 MB meeting [by Markus Schulz](#)
 - ◆ In a nutshell:
 - ◇ O(10) sites would have (part or all of) their **production** resources in "pre-prod" mode, i.e. frequently applying updates from EPEL-testing and WLCG-testing repositories
 - ◇ overlap with EGI UMD Staged Rollout participation would be good
 - ◇ those resources will be exposed to real work
 - ◇ the additional failure risks would be small
 - ◇ the benefits are for the whole infrastructure: deploy upgrades that have proved themselves (albeit at a small scale)
 - ◇ avoid ad-hoc validation tests that take significant effort to organize
 - ◆ A formalization of what was done for the EMI-2 WN validation
 - ◆ The participating sites ideally will cover:
 - ◇ All experiments
 - ◇ All services relevant per experiment
 - ◆ Let's start small with the most important use cases
 - ◇ Gain experience and adjust
 - ◆ Timeline
 - ◇ It should work by next spring
 - ◇ TF really starts in Sep?
 - ◇ Try to have resources committed by Oct
 - ◆ Sites?
 - ◆ People?

Ian is concerned about a plan that foresees to have "fragile" worker nodes in production and no rollback procedures. Maarten adds that in the past there were few incidents requiring a rollback and that rollback is usually (but not always) easy.

Ian expresses surprise about the fact that the task force was presented (and approved) at the WLCG MB earlier than in a WLCG operations meeting, as the procedure foresees to first discuss new task force proposals here. He recommends for the future to respect the scope of the various meetings.

Maria points out that the needed effort won't be small and it will need another co-chair (Maarten being one) and volunteer sites.

Helge is concerned about the fact that sites using part of their production resources for this activity might be negatively impacted in terms of availability; Maarten says that in such case sites should not be "blamed".

Alessandra F. strongly disagrees with having sites test on production resources.

Maria proposes to have a planning meeting in September to define a clear mandate and plan for the task force. The proposal is approved.

Middleware news and baseline versions

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

Highlights:

- The new baseline version for dCache is 2.2. Version 1.9.12 will reach end of support on August 31.
- perfSONAR has been added to the table
- CVMFS and StoRM versions have been changed (see the notes in the table)

Tier-1 Grid services

Storage deployment

Site	Status	Recent changes	Planned changes
CERN	CASTOR: v2.1.13-9-2 and SRM-2.11 for all instances (SRM-2.11-2 for LHCb) EOS: ALICE (EOS 0.2.37 / xrootd 3.2.8) ATLAS (EOS 0.2.38 / xrootd 3.2.8 / BeStMan2-2.2.2) CMS (EOS 0.2.29 / xrootd 3.2.7 / BeStMan2-2.2.2) LHCb (EOS 0.2.29 / xrootd 3.2.7 / BeStMan2-2.2.2)	SRM-LHCb update after repeated crashes EOSATLAS, EOSALICE updates to 0.2.38	
ASGC	CASTOR 2.1.13-9 CASTOR SRM 2.11-2 DPM 1.8.6-1 xrootd 3.2.7-1	None	None
BNL	dCache 2.2.10 (Chimera, Postgres 9 w/ hot backup) http (aria2c) and xrootd/Scalla on each pool	None	None
CNAF	StoRM 1.8.1 (Atlas, CMS, LHCb)		

FNAL	dCache 1.9.5-23 (PNFS, postgres 8 with backup, distributed SRM) httpd=2.2.3 Scalla xrootd 2.9.7/3.2.7-2.osg Oracle Lustre 1.8.6 EOS 0.2.38/xrootd 3.2.7-2.osg with Bestman 2.2.2.0.10	Upgraded EOS	FTS3, EOS 0.3, dCache 2.2 + Chimera
IN2P3	dCache 2.2.12-1 (Chimera) on SL6 core servers and 2.2.13-1 on pool nodes Postgres 9.1 xrootd 3.0.4	none	none
KISTI	xrootd v3.2.6 on SL5 for disk pools xrootd 20100510-1509_dbg on SL6 for tape pool dpm 1.8.6	xrootd upgrade on disk pools (20100510-1509_dbg -> v3.2.6)	xrootd upgrade foreseen for tape (20100510-1509_dbg -> v3.1.1) in September
KIT	dCache <ul style="list-style-type: none"> • atlasrm-fzk.gridka.de: 1.9.12-11 (Chimera) • cmssrm-fzk.gridka.de: 1.9.12-17 (Chimera) • lhbsrm-kit.gridka.de: 1.9.12-24 (Chimera) xrootd (version 20100510-1509_dbg and 3.2.6)		Upgrading all dCache instances and the PostgreSQL databases around the GridKa "firewall downtime" in CW 30 (22th-26th July)
NDGF	dCache 2.3 (Chimera) on core servers. Mix of 2.3 and 2.2 versions on pool nodes.		
NL-T1	dCache 2.2.7 (Chimera) (SURFsara), DPM 1.8.6 (NIKHEF)		
PIC	dCache head nodes (Chimera) and doors at 1.9.12-23 xrootd 3.3.1-1	head nodes to 1.9.12-23	Next upgrade to 2.2 in September
RAL	CASTOR 2.1.12-10 2.1.13-9 (tape servers) SRM 2.11-1		Upgrading all instances to CASTOR 2.1.13-9 by end of July.
TRIUMF	dCache 2.2.13(chimera), pool/door 2.2.10	Webdav door is open to ATLAS	

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		
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		During the site wide downtime 24./25.07. FTS2 machines will be reinstalled and another machine will be added.
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	None
CERN	1.8.6-1	SLC6, EMI2	Oracle 11	ATLAS, LHCb, OPS, ATLAS Xroot federations	

Other site news

Data management provider news

Experiment operations review and plans

ALICE

- CVMFS
 - ◆ The deployment campaign was started on July 4: thanks to Stefan and Guenter!
 - ◇ 21 tickets already solved/verified, 35 still open
 - ◆ A dedicated CREAM CE + WN have been set up to speed up testing of the necessary AliEn adjustments
- CERN
 - ◆ On Thu June 27 many thousands of "ghost" jobs were found keeping job slots occupied, due to aria2c Torrent processes not exiting after their 15 minutes of lifetime.
 - ◆ It is still not known what would have caused this change of behavior (no relevant changes known on the ALICE side).
 - ◆ The matter also had an impact on the routers connecting the WN subnets: their CPU usage had gone up by a lot since the evening of Fri June 21, putting the network stability at risk.
 - ◆ To mitigate the situation for the weekend, Fri June 28 late afternoon the ALICE LSF quota was reduced from 15k to 7500 jobs and a 7k cap was applied on the ALICE side as well, just to be sure.

- ◆ During the weekend we tested and deployed a patch for the ALICE job wrapper that now kills any such processes explicitly.
- ◆ Since the new release started getting used on the WN, no such lingering processes were seen any more!
- ◆ The differences with the previous release do not explain that change in behavior.
- ◆ The problem was also seen at RAL and one or two T2, while the majority of Torrent sites did not report anything amiss.
- CERN
 - ◆ Alarm ticket GGUS:95662 [↗](#) on Thu July 11 when almost all CEs refused proxies of ALICE (and other VOs) due to an Argus configuration mishap.
- CNAF
 - ◆ A plan has been developed for re-staging 2010 data (400k files) to check for corrupted files (GGUS:95073 [↗](#)) and have such cases fixed, while avoiding contention with reprocessing campaigns: thanks!
- KIT
 - ◆ Concurrent jobs cap has been kept at 2k since a week to avoid firewall overload while the local SE issues could not yet be worked on.

ATLAS

- webDAV: last meeting has been discussed the possibility of having webDAV under baseline versions. Any news?
- we would also like to take the opportunity of this meeting make a question: is there any news about WLCG accounting for MCORE? Where can we check those information? Thks

Nicolò answers that he already wrote a twiki with the WebDAV port values by storage system; the port is configurable in most implementations and there is no unique default number. Ueda wonders if WLCG should request a standard port number (but values like 443 would attract unwanted traffic). Maarten agrees that some standardisation would be desirable. Nicolò will look what numbers the sites actually use and will come with a proposal for the next meeting.

About multicore accounting, the question should best be asked to Michel.

CMS

update on various projects:

- cvmfs:
 - ◆ were advised by Jakob Blom to split the CVMFS directory by SCRAM_ARCH, done from now on
 - ◆ time schedule:
 - ◇ Fall 13: active installation at sites will be stopped, sites need to have either CVMFS or install the releases themselves through cron jobs
 - ◇ Spring 14: installation basis will be CVMFS, no cron installation, either CVMFS or CVMFS over NFS or similar
- multi-core
 - ◆ dynamic partitioning in operation (run single core jobs in multi-core pilots)
 - ◆ extending the available multi-core resources to run normal single-core production workflows to gain experience
- glideIn WMS setup
 - ◆ simplified setup to have one frontend for production and one frontend for analysis, Condor 8.0.1 has the condor_ha fix which will be used now to allow for redundancy on collector level
- opportunistic resource usage

- ◆ Parrot now uses UI via cvmfs from /cvmfs/grid.cern.ch (for now using the gLite UI because instead of the EMI UI, apart from missing init scripts the EMI UI also don't have any CA certs.)
- T1 disk/tape separation:
 - ◆ T1_UK_RAL: in operation, prototype site using for production in separated mode regularly now
 - ◆ T1_IT_CNAF and T1_ES_PIC: currently commissioning PhEDEx links for Disk endpoints
 - ◆ T1_DE_KIT: Second SRM endpoint for Disk upcoming
 - ◆ T1_FR_CCIN2P3: Possibly namespace separation in September site downtime
 - ◆ T1_US_FNAL: two independent dCache instances: New instance for Disk and Current instance (to be upgraded to Chimera + dCache 2.2 in Summer) for tape

It would be good to make the EMI UI in grid.cern.ch fully usable. Oliver G. will create a GGUS ticket to the tarball support group. UPDATE: ticket created, GGUS:96030.

LHCb

- GRIDKA : a solution seems to have improved the situation for the tape system but we still have a huge backlog so if the site could help to speed up the process we would appreciate it.

Thomas explains that the problem was that requested files were recalled from tape and copied to the staging pools, but the last hop to the online space was not performed. Now all files are automatically copied to the read pools.

News from EGI Operations

Peter gives a report, these being the highlights:

- Deployed middleware is being checked for SHA-2 compliance by SAM/Nagios (MIDMON instance); in a couple of days sites will start receiving tickets. There is no hard deadline yet to upgrade to SHA-2-compliant versions.
- At the EGI TF 2013 in Madrid (September 16-20) there will be trainings for site admins, including on how to properly publish and debug GLUE2 site information in BDII.
- There is a discussion about the possibility to include frontier-squid in UMD, which might be convenient for sites; the developer (Dave) is open to the possibility. However this implies a verification and staged rollout process and early adopters. It is important to understand how many sites would find it beneficial.
- It has been observed that only ~10 sites publish their squid servers in GOCDB, while there are ~200 sites in WLCG with a frontier-squid server.

Maarten thinks that it is quite possible to have a vast majority of services running SHA-2-compliant versions by the end of the year. OSG is much more advanced than EGI and there is no need to followup from operations.

Nicolò and Peter add that the SHA-2 compatibility will be backported to dCache 2.2 and this version will also be available in UMD.

Alessandro reminds that in fact sites are requested to publish their frontier-squid servers in GOCDB/OIM.

About having frontier-squid in UMD, Andrea argues that it should be irrelevant for the experiments, so the question is to the sites. He suggests that those sites asking for it should volunteer as early adopters. Peter concludes that EGI will collect more feedback from the sites, but clearly it won't be done for the sake of just a handful of sites, if most sites would anyway use other installation methods.

Task Force reports

SL6

- T1s Done: 7/15 (Alice 4/9, Atlas 5/12, CMS 3/9, LHCb 4/8)
 - ◆ +1 since last update
 - ◆ All T1s now have a migration plan
 - ◆ TRIUMF has gone online last week with a fraction of resources and will complete the migration by 22/7/2013
- T2s Done: 35/129 (Alice 7/39, Atlas 17/89, CMS 18/65, LHCb 9/45)
 - ◆ +7 since last update
 - ◆ Only 36 remain without any plan or testing going on.
- EMI-3 testing
 - ◆ voms-proxy-info: had another problem: java client memory request wasn't limited and clashed with sites setting vmem limits on the WNs. It affected atlas jobs. The current version is affecting also WLCG VOBOXES * Memory limit have now been set in a new test version which has solved the atlas problems at 2 UK sites. One site is now online with both production and analysis queues since Monday. However it is clear these new VOMS clients need more testing as they got 3 tickets in a month for different problems. CMS and VOBOXES should let me know if the version with memory limits set works for them. Tickets are GGUS: 94878 [↗](#), 95574 [↗](#), 95798 [↗](#)
- Atlas had a number of unexpected problems that slowed the sites migration down but they are all solved now.

Maarten clarifies that for the WLCG VOBOX the problem is with voms-proxy-init: Myproxy does not work with VOMS proxies generated by the new client.

SHA-2 migration

- EGI have added SHA-2 tests [↗](#) to the middleware monitoring service ("midmon"), currently checking the following services for SHA-2 readiness:
 - ◆ CREAM-CE (eu.egi.sec.CREAMCE-SHA-2) - 176 instances in warning
 - ◆ StoRM (eu.egi.sec.StoRM-SHA-2) - 46 instances in warning
 - ◆ VOMS (eu.egi.sec.VOMS-SHA-2) - 38 instances in warning
 - ◆ WMS (eu.egi.sec.WMS-SHA-2) - 41 instances in warning
- SHA-2 support middleware baseline [↗](#)
- dCache
 - ◆ version 2.6.5 released July 16 provides SHA-2 support
 - ◇ see the release notes [↗](#)
 - ◆ on July 23 the last 2.2.x **without** SHA-2 support will be released
 - ◆ on July 30 the first 2.2.x **with** SHA-2 support will be released

FTS-3

A pre-production CERN FTS3 service has been deployed by CERN-IT-PES for testing: CMS reconfigured the CERN PhEDEx instance to use fts3.cern.ch for LoadTest transfers to CERN. More news soon. Also LHCb has started using the new CERN FTS3 for 5 sites.

Starting last week (10 July) ATLAS started using FTS3 RAL production instance for all the activities for those sites: UKI-NORTHGRID-LANCS-HEP, UKI-NORTHGRID-MAN-HEP, UKI-SCOTGRID-ECDF, UKI-SOUTHGRID-RALPP, WEIZMANN-LCG2, TECHNION-HEP, IL-TAU-HEP, RU-Protvino-IHEP. Progress can be tracked <https://savannah.cern.ch/bugs/index.php?102004> [↗](#) From this link you may find all the ATLAS DDMEndpoints served by FTS3 instances (CERN pilot for now, plus RAL)

http://atlas-agis.cern.ch/agis/ddm_endpoint/table_view/?&state=ACTIVE&fts=FTS3 No problem observed up to now, ATLAS plan is to keep adding sites for all production activities. The whole UK ATLAS cloud is also served by RAL FTS3 for Functional Test activity. CMS is sending debug transfers through RAL FTS3 instance to 5 sites. Plan is to keep on adding sites for debug transfers.

gLExec

- Deployment tracking page
- 25 tickets solved and verified, 68 still open, many on hold until late summer/autumn

perfSONAR

- wlcg mesh creates big log files up to the point of exhausting the machine disk space. The problem is under investigation by the TF in cooperation with the perfsonar-PS developers. The problem is caused by the pinger producing an abnormal number of messages. A second issue is always with the pinger actually not performing any tests despite the mesh being properly populated in the list of tests on the latency host. There is already a fix that might solve the problem which is under test at AGLT2. More info in the next few days. The fix might be released as a minor perfsonar update 3.3.1 or adding to the current repo this is also under discussion with the developers. Sites should avoid using the wlcg mesh test until this is fixed.

AOB

Action list

1. Build a list by experiment of the Tier-2's that need to upgrade dCache to 2.2
 - ◆ **done** for ATLAS: list updated on 17-05-2013
 - ◆ **done** for CMS: list updated on 20-06-2013
 - ◆ not applicable to LHCb, nor to ALICE
 - ◆ Maarten: EGI and OSG will track sites that need to upgrade away from unsupported releases.
2. Inform sites that they need to install the latest Frontier/squid RPM by May at the latest (**done** for CMS and ATLAS, status monitored[?])
3. Maarten will look into SHA-2 testing by the experiments when experts can obtain VOMS proxies for their VO.
 - ◆ in progress, see SHA-2 readiness testing
4. 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.
5. Investigate how to separate Disk and Tape services in GOCDB
 - ◆ proposal submitted via GGUS:93966[?]
6. Agree with IT-PES on a clear timeline to migrate OPS and the LHC VOs from VOMRS to VOMS-Admin
7. For the experiments to give feedback on the machine/job information specifications (**done** as now managed by task force)
8. Experiments interested in using WebDAV should contact ATLAS to organise a common discussion
9. Add KISTI to the list of Tier-1 sites in the Grid Services report
 - ◆ **done**
10. Contact the storage system developers to find out which are the default/recommended ports for WebDAV
 - ◆ **in progress** Recommendations for WebDAV deployment
11. Circulate the instructions to enable the xrootd monitoring on DPM
 - ◆ **done** <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Xroot/Setup#VOcentralmonitoring>[?]

-- AndreaSciaba - 17-Jun-2013

This topic: LCG > WLCGOpsMinutes130718

Topic revision: r33 - 2013-07-23 - AndreaSciaba



Copyright &© 2008-2022 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](#)