

Table of Contents

WLCG Operations Coordination Minutes - March 20, 2014.....	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.....	3
Oracle deployment.....	3
Other site news.....	4
Wigner job efficiency state of affairs.....	4
Data management provider news.....	5
Experiments operations review and Plans.....	5
ALICE.....	5
ATLAS.....	5
CMS.....	6
LHCb.....	6
Ongoing Task Forces and Working Groups.....	6
Tracking tools evolution TF.....	7
FTS3 Deployment TF.....	7
gLExec deployment TF.....	7
Machine/Job Features.....	7
Middleware readiness WG.....	7
Multicore deployment.....	7
perfSONAR deployment TF.....	7
SHA-2 Migration TF.....	8
WMS decommissioning TF.....	8
xrootd deployment TF.....	8
IPv6 validation and deployment TF.....	8
HTTP proxy discovery TF.....	8
Action list.....	8
AOB.....	9

WLCG Operations Coordination Minutes - March 20, 2014

Agenda

- <https://indico.cern.ch/event/308833/>

Attendance

- Andrea Sciaba (chair), Nicolo Magini (secretary)
- local: Maria Alandes, Simone Campana, Marcin Blaszczyk, Maria Dimou, Oliver Keeble, Stefan Roiser, Alesandro Di Girolamo, Felix Lee, Daniele Spiga, Maarten Litmaath, Pablo Saiz
- remote: Valery Mitsyn, Yury Lazin, Maite Barroso, Daniele Bonacorsi, Thomas Hartmann, Christoph Wissing, Massimo Sgaravatto, Antonio Perez, Alessandra Forti, Alessandro Cavalli, Frederique Chollet, Pepe Flix, Rob Quick, Michel Jouvin

News

- Simone Campana is the new ATLAS Distributed Computing coordinator and Andrea Sciabà and Maria Alandes are the new WLCG operations officers
- Alessandra announced that the last T2 migrated to SL6 (in this case, CentOS 6.x). The only remaining site is CERN, which has moved 62% of the resources to SL6
- Alessandra started a web page for batch system comparison as decided at the latest pre-GDB meeting
- GOCDB now supports custom tags that can be attached to sites and services. This can be extremely useful for operations (e.g. they can be used to expose the VOs supported by a service) but it might require some previous agreement on the names of tags to introduce
- It has been proposed to move the planning meeting from April 3 to April 17 to prevent a conflict with the HEP software collaboration meeting that will take place on April 3
- Sites are reminded that CVMFS 2.0 is reaching end-of-life at the end of March and they should upgrade their CVMFS clients as soon as possible, as it's a prerequisite for the upgrade of the Stratum0 repositories, planned for June
- Experiments are reminded that the plan is to migrate to the GFAL2/FTS3 clients and decommission the GFAL/FTS2 clients during the summer
- **The April 1 deadline for deploying the latest version of perfSONAR and opening the firewall to the IPs needed for the central monitoring is very close and many sites are not OK: please take action!**
- Proposed dates of future meetings:

Date	Type	Notes
3/4	coordination	
17/4	planning	
8/5	coordination	shifted due to May 1
22/5	coordination	shifted due to May 1
5/6	coordination	
19/6	coordination	
7-9/7	workshop	in Barcelona
24/7	coordination	shifted due to workshop

- Alessandro asks if OIM also supports custom tags like GOCDB. Rob answers that OIM already has VO support tags, other usecases to be identified.

Middleware news and baseline versions

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

- perfSONAR baseline increased to 3.3.2 as requested by TF

Tier-1 Grid services

Storage deployment

Site	Status	Recent changes	Planned changes
CERN	CASTOR: v2.1.14-11 and SRM-2.11-2 on ATLAS, ALICE, CMS (LHCB running 2.1.14-5 will be updated next week) EOS: ALICE (EOS 0.3.4 / xrootd 3.3.4) ATLAS (EOS 0.3.8 / xrootd 3.3.4 / BeStMan2-2.3.0) CMS (EOS 0.3.7 / xrootd 3.3.4 / BeStMan2-2.3.0) LHCb (EOS 0.3.3 / xrootd 3.3.4 / BeStMan2-2.3.0 (OSG pre-release))	CASTOR (except LHCb) upgraded	LHCb CASTOR upgrade next week
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 9.3 w/ hot backup) http (aria2c) and xrootd/Scalla on each pool	None	None
CNAF	StoRM 1.11.3 emi3 (ATLAS, LHCb, CMS)	also CMS updated to 1.11.3	
FNAL	dCache 2.2 (Chimera, postgres 9) for disk instance; dCache 1.9.5-23 (PNFS, postgres 8 with backup, distributed SRM) for tape instance; httpd=2.2.3 Scalla xrootd 2.9.7/3.2.7.slc EOS 0.3.15-1/xrootd 3.3.6-1.slc5 with Bestman 2.3.0.16		
IN2P3	dCache 2.6.18-1 (Chimera) on SL6 core servers and pool nodes Postgres 9.2 xrootd 3.3.4 (Alice T1), xrootd 3.3.4 (Alice T2)		
JINR-T1	dCache <ul style="list-style-type: none"> • srm-cms.jinr-t1.ru: 2.6.23 • srm-cms-mss.jinr-t1.ru: 2.2.24 with Enstore xrootd federation host for CMS: 3.3.6		
KISTI	xrootd v3.2.6 on SL5 for disk pools xrootd 20100510-1509_dbg on SL6 for tape pool dpm 1.8.7-3		
KIT	dCache <ul style="list-style-type: none"> • atlassrm-fzk.gridka.de: 2.6.21-1 • cmssrm-kit.gridka.de: 2.6.17-1 • lhcb-srm-kit.gridka.de: 2.6.17-1 	None	None

	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 redirector 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 2.2.23-1 xrootd door to VO servers (3.3.4)	None	None
RAL	CASTOR 2.1.13-9 2.1.14-5 (tape servers) SRM 2.11-1		
RRC-KI-T1	dCache 2.2.24 + Enstore (ATLAS) dCache 2.6.22 (LHCb) xrootd - EOS 0.3.19 (Alice)		
TRIUMF	dCache 2.6.21		

- New T1 RRC-KI-T1 in Russia supporting ALICE, ATLAS and LHCb added to the table

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		
JINR-T1	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	None	None
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	1.8.7-3	SLC6, EPEL	Oracle 11	ATLAS, LHCb, OPS, ATLAS Xroot federations	

Oracle deployment

- Note: only Oracle instances with a direct impact on offline computing activities of LHC experiments are tracked here

- Note: an explicit entry for specific instances is needed only during upgrades, listing affected services. Otherwise sites may list a single entry.

Site	Instances	Current Version	WLCG services	Upgrade plans
CERN	CMSR	11.2.0.4	CMS computing services	Done on Feb 27th
CERN	CASTOR Nameserver	11.2.0.4	CASTOR for LHC experiments	Done on Mar 04th
CERN	CASTOR Public	11.2.0.4	CASTOR for LHC experiments	Done on Mar 06th
CERN	CASTOR Alicestg, Atlasstg, Cmsstg	11.2.0.4	CASTOR for LHC experiments	Done: 10-14th March
CERN	LCGR	11.2.0.4	All other grid services (including e.g. Dashboard, FTS)	Done: 18th March
CERN	CASTOR LHCbstg	11.2.0.3	CASTOR for LHC experiments	upgrade to 11.2.0.4 planned for 25th March
CERN	LHCBR	11.2.0.3	LHCb LFC, LHCb Dirac bookkeeping	upgrade to 12.1.0.1 planned for 24th of March
CERN	ATLR, ADCR	11.2.0.3	ATLAS conditions, ATLAS computing services	TBA: upgrade to 11.2.0.4
CERN	HR DB	11.2.0.3	VOMRS	upgrade to 11.2.0.4 planned for 14th April
CERN	CMSONR_ADG	11.2.0.3	CMS conditions (through Frontier)	TBA: upgrade to 11.2.0.4 (tentatively May)
BNL		11.2.0.3	ATLAS LFC, ATLAS conditions	TBA: upgrade to 11.2.0.4 (tentatively June)
RAL		11.2.0.3	ATLAS conditions	TBA: upgrade to 11.2.0.4 (tentatively June)
IN2P3		11.2.0.3	ATLAS conditions	TBA: upgrade to 11.2.0.4 (tentatively March?)
TRIUMF	TRAC	11.2.0.4	ATLAS conditions	Done

- Need confirmation from IN2P3 about Oracle upgrade in March
- LFC and DIRAC bookkeeping were successfully validated on Oracle12

Other site news

Wigner job efficiency state of affairs

- some significant differences in job efficiency (CPU/wall-clock) are observed when the Wigner and Meyrin computer centers are compared
- the effects differ per experiment and per workflow
 - ◆ analysis jobs typically are affected, up to ~20%
 - ◆ simulation typically is not
- the investigation is complex because of multiple factors:
 - ◆ network latency and performance between Wigner and Meyrin
 - ◇ in particular for accessing EOS at CERN (for the time being the LHC experiments data are only in Meyrin)
 - ◇ mitigations are being looked into
 - ◆ Wigner and Meyrin have different CPU type mixes (AMD vs. Intel)
 - ◆ Wigner nodes are all virtualized, while Meyrin also has physical nodes
 - ◆ Wigner nodes run only SLC6, Meyrin has SLC5 and SLC6
 - ◆ for each experiment both centers receive job mixes that may not be easy to disentangle

- ◊ we cannot simply reserve a lot of capacity for dedicated tests, because the resources are needed for real work!
- ◆ the investigations require a number of experts to be involved, and such people are generally quite busy already
- to be continued...

- capacity numbers determined on March 11:
 - ◆ Wigner: 7976 non-dedicated slots, 57613 HS06
 - ◆ Meyrin: 32236 non-dedicated slots, 288042 HS06
 - ◆ SLC5: 118897 HS06
 - ◆ SLC6: 226781 HS06

- Cristoph asks how to target intentionally Meyrin vs Wigner in job submission. Maarten answers that it is not easy, but possible tweaking the experiment frameworks. As it is not the submission mode intended by IT, the message is that there is no reason to make difference for real production, only for debugging.

Data management provider news

Experiments operations review and Plans

ALICE

- T1-T2 workshop [↗](#), March 3-7, Tsukuba, Japan
 - ◆ admins please have a look at the relevant presentations, thanks!
- The productions and analysis for Quark Matter 2014 [↗](#) are progressing at steady pace and according to plan.
- CERN
 - ◆ Wigner job efficiencies
 - ◊ no change
 - ◊ next meeting Apr 4

ATLAS

- Rucio: migration from LFC to RFC is progressing without problems, as of today we have migrated one entire cloud plus few more sites. In the next weeks we will increase the speed of how many sites at once we migrate.
- JEDI for analysis: few users tested without problems. In the next weeks few fixes/improvements in the monitoring will be done and we can then go in full production.
- Activity as usual: MC production run 30% MCORE.
- slow transfers have been observed between various SRC and DST. 3 main different issues: UKI-SOUTHGRID-CAM-HEP -> BNL, NIKHEF->TRIUMF, NDGF->TAIWAN. For CAM-BNL transfers with FTS3 CERN were 0.4MB/s, while with BNL FTS2 20MB/s. But then we noticed that lcg-cp with >3 streams, srmcp, and also FTS2 RAL managed transfers have the same performances. We also noticed that a very similar problem was observed already in the past http://bourricot.cern.ch/dq2/ftsmon/multi_spacetoken_view/UKI-SOUTHGRID-CAM-HEP/BNL-OSG2/2013.
- regarding the Geneva-Wigner measurements for ATLAS: the impact on production activity is not really visible, on analysis we observe 20% efficiency degradation. The overall total impact on the ATLAS resources is of the order of few % (having ATLAS approx 90-10 prod vs analysis share), but it's highly visible because it's affecting user jobs.
 - ◆ we are now in a situation in which we are able to study and measure the issues and their evolution with time

- ◆ network experts have been included in the loop
 - ◆ mitigation measures have been taken by some experiment (e.g. for ATLAS activation of TTReeCache when possible for analysis activity is ongoing)
 - ◆ for the future we also believe that once the data will be spread between Wigner and Meyrin the geolocation of the client should mitigate even further the observed issues because the "closer" copy of the file will be read.
- Comment that the "slow" links seem to be only transcontinental

CMS

- FTS2 decommissioning at CERN by August 1st?
 - ◆ CMS will come to an answer next week (during C&O workshop)
 - FTS3
 - ◆ Sites should switch their Phedex "Debug" agents to use RAL FTS3
 - Brief production and processing overview
 - ◆ HeavyIon rereco pass lunched last week
 - ◇ Quite memory (3-4GB) and cpu (96h jobs) demanding
 - ◇ Also using HLT-Cloud
 - ◆ Phase II upgrade MC
 - ◆ Soon 13 TeV MC digi/reco
 - CPU efficiency for jobs running in Wigner center
 - ◆ One reported case with rather bad CPU efficiency
 - ◇ Quite special work flow
 - ◇ Limited statistics
 - ◆ More general analysis ongoing
 - CVMFS Migration
 - ◆ Finalizing to use CMSSW via CVMFS at CERN
 - ◇ Expect to switch by April 1st
 - Multicore processing
 - ◆ Switching a large fraction of FNAL resources
 - ◆ Continue testing at other sites, primarily Tier-1s
 - CMS Spring Computing & Offline week at CERN * March 24-28
- Maria Dimou reminds to inform CMS users about deployment of new features for CMS in GGUS on Wed 26th. The target for disabling the Savannah-to-GGUS bridge functionality in GGUS is one month later. Cristoph says that CMS Computing Operations workflows will be moved from Savannah to GGUS piecewise.

LHCb

- Spring '14 Incremental stripping campaign
 - ◆ 2/3 of the data have been processed, estimating another 2 weeks of processing
 - ◆ Several sites have finished staging and processing of data
- WMS decommissioning campaign for remaining few small sites resurrected
 - ◆ Some problems found with parsing of BDII information by Dirac
- FTS3 clients compilation in AA context currently on hold b/c of C++98/11 libraries (AA / middleware)

Ongoing Task Forces and Working Groups

Tracking tools evolution TF

- GGUS release on the 26th of March. The release will bring the possibility to notify multiple sites with one ticket, Shibboleth support, and implement several CMS specific requests. The service will be in downtime from 7:00 to 10:00 UTC

FTS3 Deployment TF

- Understanding slow transfers as reported by ATLAS. We have to study the behaviour of the optimizer in these "strange" situations.
- The FTS3 server was not using the correct user proxy to transfer their files if the user has used several proxies (each one with different role) for submitting FTS jobs to the server. Developers have implemented the fix and applied it in the FTS3 pilot server for testing.

gLExec deployment TF

- 79 tickets closed and verified, 16 still open
- Deployment tracking page

Machine/Job Features

- patch release for the mjf.py client made available for CERN batch nodes. Deployment by PES

Middleware readiness WG

- 3rd meeting of the WG held on March 18. Watch the Agenda link for the minutes [↗](#) which are being prepared.
- Next (4th) meeting will take place on Thursday May 15 at 10:30 CEST. Items in that agenda will be:
 - ◆ The different Readiness Verification approaches across experiment VOs.
 - ◆ The refinement of the WLCG MW Officer role.
 - ◆ Issues of exact versions run at the *Volunteer* sites.

Multicore deployment

- NTR

perfSONAR deployment TF

- We have added 3 new sites in the LHCONe backbone and a new site in Finland is "in-progress", waiting on a decision about which mesh to put them in.
- In terms of "quality" we have seen some definite improvements in the CA, FR, USATLAS, USCMS and UK clouds.
- We still have quite a few issues with getting tests to run and return the needed metrics. This is caused by firewalls, down services or sites not yet using the mesh-config. We have a target of April 1, 2014 to have all sites instrumented.
- Apologies to the sites which have been waiting for an answer in the tickets assigned to them
- Less than 10 sites have not yet installed perfSONAR; but many sites are still running old releases (especially Tier-1s)

SHA-2 Migration TF

- SHA-2 user certificates are being used by all 4 experiments
- campaign to get our future VOMS servers recognized across WLCG has been launched on Mon March 17
 - ◆ EGI broadcast [↗](#)
 - ◆ OSG operations team is involved
- SAM tests could be used to discover if new VOMS servers are recognized, though not optimal
- The action about VOMS servers is not SHA-2 specific, so it can be decoupled and moved to the action list. Since there are no remaining actions concerning SHA-2, it was agreed to review the TF status at the planning meeting on April 17th and possibly declare it concluded.

WMS decommissioning TF

- CERN WMS instances for experiments will start getting drained as of April 1
- SAM instances have their own timeline
- Andrea comments that the timeline for the SAM WMS decommissioning is the validation of the new SAM probes with direct CREAM and CondorG submission, which is on schedule for June

xrootd deployment TF

- NR

IPv6 validation and deployment TF

- ATLAS first goal to be able to run Panda Tasks on IPv6 WN.
 - ◆ Panda Client uses http (curl) so software should work.
 - ◆ Development Panda instances at CERN is being made dual stack.
 - ◆ Development Pilot Factory instance at CERN is being made dual stack
- Frontier/Squid test box has also been made dual stack.

HTTP proxy discovery TF

* NR

Action list

1. Investigate how to separate Disk and Tape services in GOCDB
 - ◆ proposal submitted via GGUS:93966 [↗](#)
 - ◆ **in progress** - ticket updated, current solution to be validated.
 - ◇ Some of the T1 sites are adding `SRM.nearline` entries as desired.
 - ◇ Downtime declaration tests to be done.
 - ◇ Experiences to be reported in the ticket.
2. Agree with IT-PES on a clear timeline to migrate OPS and the LHC VOs from VOMRS to VOMS-Admin
 - ◆ **in progress**
 - ◇ voms-admin test cluster is available
 - ◇ experiment VO managers have given feedback
 - ◇ VOMS service managers are in contact with developers to get major issue(s) fixed
 - ◇ <https://twiki.cern.ch/twiki/bin/view/GridServices/VomsAdminTestInstance>

- About the GOCDB action:
 - ◆ Stefan is coordinating for LHCb Tier-1s, though DIRAC is not ready yet.
 - ◆ For ATLAS and CMS, the statement is still:
 - ◇ no problem if Tier-1s declare new "SRM.nearline" services, but not needed yet. Will be evaluated in the future.
 - ◇ Tier-1 "SRM" services should not be in "OUTAGE" if only tape is affected. Declaring "AT RISK" downtime is fine.
- About VOMS-Admin, Maarten comments that several fixes were applied reducing the discrepancies between HR DB and VOMS DB, but not yet fully fixed. No significant problems reported by VO Managers with GUI or API so far.

AOB

-- Nicolomagini - 18 Mar 2014

This topic: LCG > WLCGOpsMinutes140320

Topic revision: r36 - 2014-03-31 - ValeryMitsyn



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