

Table of Contents

WLCG Operations Coordination Minutes - March 6, 2014.....	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.....	4
Oracle deployment.....	4
Other site news.....	5
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.....	7
FTS3 Deployment TF.....	7
gLExec deployment TF.....	7
Machine/Job Features.....	7
Middleware readiness WG.....	7
Multicore deployment.....	7
perfSONAR deployment TF.....	8
SHA-2 Migration TF.....	8
Tracking tools evolution TF.....	9
WMS decommissioning TF.....	9
xrootd deployment TF.....	9
Action list.....	9
AOB.....	9

WLCG Operations Coordination Minutes - March 6, 2014

Agenda

- <https://indico.cern.ch/event/280062/>

Attendance

- Alessandra Forti (chair), Nicolo' Magini (secretary)
- Local: Andrea Sciaba', Michail Salichos, Stefan Roiser, Marcin Blaszczyk, Maria Alandes, Simone Campana, Alessandro Di Girolamo, Maria Dimou
- Remote: Javier Sanchez, Thomas Hartmann, Yury Lazin, Alessandra Doria, Antonio Perez Calero Yzquierdo, Christoph Wissing, Di Qing, Diego Gomes, Gareth Smith, Maite Barroso Lopez, Frederique Chollet, Alessandro Cavalli, Shawn Mc Kee, Peter Gronbech

News

- Simone was nominated ATLAS Distributed Computing coordinator and will step down as chair of WLCG Operations Coordination. Waiting for official communication on who will take over his duties.
- The schedule of upcoming meetings will be circulated after the meeting. Dates in May are shifted by one week to accommodate holidays and the workshop
- Reminder about pre-GDB on batch systems next week in Bologna, attendance from sites is encouraged. One of the main topics of discussion will be MAUI/torque, since MAUI is unsupported. Multicore support will also be discussed
- Alessandro comments about overlap between multicore Task Force and pre-GDB on batch systems, which makes difficult for people to follow all discussions. Acknowledged, though multicore will not be the only topic at the pre-GDB

Middleware news and baseline versions

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

- Baseline highlights: WMS fix for 512-bit keys, already applied at CERN. Maite comments that WMS update are still applied as needed since decommissioning deadline has not yet been met, and there is agreement for support for SAM

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 0.3.8 / xrootd 3.3.4 /		ongoing: CASTOR DB hardware migration+updates to ORACLE11.2.0.4 (downtime), combined with roll-out of CASTOR 2.1.14-11

Site	Status	Recent changes	Planned changes
	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))		
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) StoRM 1.11.2 emi3 (CMS)		
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	Moved disk instance into production with all pools	Begin upgrade process for tape instance to dCache 2.2
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)	transition to xrootd 3.3.4 for ALICE. Issues on T1 instance under investigation.	
JINR-T1	dCache <ul style="list-style-type: none"> • srm-cms.jinr-t1.ru: 2.6.21 • srm-cms-mss.jinr-t1.ru: 2.2.24 with Enstore xrootd federation host for CMS: 3.3.3		
KISTI	xrootd v3.2.6 on SL5 for disk pools xrootd 20100510-1509_dbg on SL6 for tape pool dpm 1.8.7-3	None	None
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 xrootd	<ul style="list-style-type: none"> • updated atlassrm-fzk.gridka.de to 2.6.21 • updated FAX pool monitoring plugins to 5.0.5-0 	

Site	Status	Recent changes	Planned changes
	<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 2.2.17-1 xrootd door to VO servers (3.3.4)		
RAL	CASTOR 2.1.13-9 2.1.14-5 (tape servers) SRM 2.11-1		Ready for 2.1.14 upgrade, date TBA. Probably non-T1 instances by end of March, T1 in April
TRIUMF	dCache 2.6.21	updated to 2.6.21	

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		
RAL	2.2.8 - transfer-fts-3.7.12-1		
TRIUMF	2.2.8 - transfer-fts-3.7.12-1		

- **Note on FTS2:** As FTS 3 is deployed in production and fully functional, CERN would like to propose a deadline to switch FTS 2 off on the 1st of August. This is because following the quattor deadline at CERN (October 2014), we would need at least 2 months so that FTS2 can be migrated to openstack, SLC6 and puppet, which is clearly something we would like to avoid. Current status:
 - ◆ LHCb - completely migrated
 - ◆ ATLAS - Well advanced , final stages, August 1st is realistic?
 - ◆ CMS - Discussion within CMS has not started.
- Cristoph comments that discussion in CMS has started
- Alessandro confirms that August 1st is OK for ATLAS, also for T1s

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	11.2.0.3	CASTOR for LHC experiments	Upgrade planned: 10-14th March
CERN	CASTOR LHCbstg	11.2.0.3	CASTOR for LHC experiments	Upgrade planned: 25th March
CERN	LHCBR	11.2.0.3	LHCb LFC, LHCb Dirac bookkeeping	TBA: upgrade to 12.1.0.1
CERN	ATLR, ADCR	11.2.0.3	ATLAS conditions, ATLAS computing services	TBA: upgrade to 11.2.0.4
CERN	LCGR	11.2.0.3	All other grid services (including e.g. Dashboard, FTS)	TBA: upgrade to 11.2.0.4 (tentatively 18th March)
CERN	HR DB	11.2.0.3	VOMRS	TBA: upgrade to 11.2.0.4 (tentatively 14th April)
CERN	CMSNR_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 middle of March)
TRIUMF	TRAC	11.2.0.4	ATLAS conditions	Done

- Marcin reports that tests of LHCBR with Oracle12 are ongoing, no issues with functionality so far, and good performance on new hardware. Tests of LFC server on Oracle12 with Oracle11 client are also ongoing, no issues seen so far.
- Marcin reports that the T1 upgrades are related to migration to golden gate (replacing streams)
- Maria comments that the tentative date of April 14th for the HR DB upgrade is just before holidays. Acknowledged, but the intervention is low risk

- Downtime of LCGR on March 18th is expected to last 2 hours. All experiments confirm that they have no problem with the date.
- Nicolo' asks if Oracle deployments at T1s for FTS2 should also be tracked given the upcoming decommissioning. No comment from the audience.

Other site news

Data management provider news

- dCache is going to extend the security support for 2.2 until Enstore and dCache 2.6 are properly integrated. This will happen by summer.

Experiments operations review and Plans

ALICE

- T1-T2 workshop [📅](#), March 3-7, Tsukuba, Japan

ATLAS

- moved all the DDM traffic to CERN FTS3 instance, due to instabilities to the virtualized infrastructure of RAL. This is just working fine. Next week we plan to mix the load, if RAL agrees.
- we are in the middle of a disk crisis, many of the Tier1s are almost full of primary data (which can't be deleted automatically). we are working to understand which kind of production generated so many data, and if a new policy (we usually keep one copy of primary of AOD,ESD, DESD on Tier1s) is conceivable.
- JEDI is under testing now. OK for HammerCloud and a small subset of users (4), we are now in the process of increasing the number of users. No problem up to now.
- JEM activated (Job Evolution Monitor) for all the production resources.
- Rucio migration (Rucio as file catalog instead of LFC) in progress. First site (LAPP) was migrated, without (major) problems.
 - ◆ we verified that the latest DQ2 clients 2.5.0 are ok everywhere. Switching just now the production CVMFS DQ2 latest link.
 - ◆ organizing the next few sites to be volunteered for the migration: at least one Tier2s from US and then a Tier1. The operation is centrally managed, and supposed to be fully transparent. If sites have in their PandaQueues allowFAX=True then we believe that we can also avoid set the site to test for the few hours needed for the migration. We are in the process of testing this.
- about Federated Data Access - from Feb 2014 ATLAS S&C Week ADC Operations session:
 - ◆ it was agreed as policy that T1s and T2Ds are to offer xrootd access to their storage, where the storage technology allows it. ADC furthermore asks and encourages sites not yet in the FAX federation to take the modest additional step beyond supporting xrootd of joining FAX. If there are technical issues, then please let ADC know.
 - ◇ We intend to demonstrate WAN data access at scale (<~10% of data access) in DC14, utilizing the technology available today: xrootd, FAX
 - Consequently, timescale for installation is in time for pre-DC14 testing
 - ◇ We intend to explore and possibly utilize HTTP as technology for federating storages and enabling WAN data access
 - Compare xrootd, http for WAN access during 2014
 - ◇ Also will put HTTP in production (e.g. downloads/dq2-get) sooner as they solve long standing issues impacting users (Does ATLAS data have to be ATLAS-only read protected?? -- discussions to be done on it)

◊ Therefore we ask sites to enable HTTP access via WebDAV on same timescale, i.e. by DC14

- Alessandro confirms that ATLAS is indeed asking all sites to enable HTTP/WebDAV permanently (even after completion of Rucio renaming).

CMS

- Current production and processing overview
 - ◆ Heavy Ion RERECO pass
 - ◆ Phase II upgrade MC
 - ◆ Soon starting 13 TeV MC DIGI/RECO
- FTS
 - ◆ FTS3 was unstable at RAL
 - ◆ Need to find a WLCG wide strategy
 - ◆ FTS2 decommissioning at CERN by Aug 1st
 - ◊ Not fully discussed in CMS yet
 - ◊ Also depends on FTS3 strategy
- Reduction of daily WLCG calls during data taking?
 - ◆ No final answer from CMS yet
- Access to high memory resources
 - ◆ Got in contact with various sites via tickets how to access them
- Multi-core
 - ◆ Want to use at least one T1 in production still in March
 - ◆ Interested sites should contact us
 - ◆ Accounting issues being discussed in the multi-core TF
- SAM submission via condor_g
 - ◆ CMS still very interested
 - ◆ Status?
- Alessandra asks CMS which T1 should be used for multi-core testing. Cristoph replies that CMS is in contact with KIT and RAL to continue testing multi-core submission, while PIC does not support larger scale activities due to accounting issues.
- Andrea confirms that the SAM team is testing condor_g probes

LHCb

- Operations
 - ◆ 28 GGUS tickets in the last 2 weeks
 - ◊ 16 tickets on pilots aborting or problems with CEs
 - ◊ 7 tickets related to software distribution
 - 2 problems with Squids, 3 problems with CVMFS clients (sites running outdated versions), 1 ticket on /cvmfs/grid.cern.ch CAs not in sync with the afs area, solved with PES
 - ◆ LHCb 2014 spring incremental stripping in full swing, 1/4 of the data has been processed.
 - ◊ Statistics available
 - <http://lhcbproject.web.cern.ch/lhcbproject/Reprocessing/stats-inc-stripping-spring14.html>
- Infrastructure

- ◆ FTS2 decommissioning
 - ◇ LHCb has fully replaced FTS2 by FTS3, therefore decommissioning is fine
- ◆ Campaign to separate disk and tape endpoints in GOCDB (see also GGUS:93966 [↗](#)).
 - ◇ Asked all LHCb supporting T1s to add "SRM.nearline" to reflect the tape endpoint. 4 sites so far have implemented this. Implementation on Dirac not yet completed -> sites who have introduced the new endpoint please put downtimes for both SRM and SRM.nearline in case of tape outage until further notice.
- Gareth asks if the "SRM.nearline" endpoints should be declared as "testing" in GOCDB, Stefan answers that it doesn't matter yet.
- ATLAS and CMS are not yet ready to make use of the downtime declarations for "SRM.nearline". However, they also confirm that they have no issue if sites declare a downtime in "SRM.nearline" for tape, as long as "SRM" is not in downtime when the disk is up.

Ongoing Task Forces and Working Groups

FTS3 Deployment TF

- Discussed with experiment DM developers how to integrate multiple FTS3 servers with experiment frameworks
- Alessandro reminds that the common strategy was also discussed with CMS, with Tony Wildish present as PhEDEx developer.

gLExec deployment TF

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

Machine/Job Features

- NTR

Middleware readiness WG

- The doodle [↗](#) showed our next (3rd) meeting is on **Tuesday 2014/03/18 @ 14:30h CET** at CERN in room 513-R-068 with audioconf. Agenda is now http://indico.cern.ch/event/MW-Readiness_3 [↗](#)
- The twiki <https://twiki.cern.ch/twiki/bin/view/LCG/MiddlewareReadinessArchive> is up-to-date.
- It contains the General Guidelines for the Readiness Verification Procedure across VOs and Products are in up for discussion at the 2014/03/18 WG meeting.

Multicore deployment

- Mini workshops on batch systems in terms of:
 - ◆ functionalities useful for multicore scheduling
 - ◆ experience so far (only ATLAS multicore jobs)
- Plan:
 - ◆ Done: HTCondor at T1_RAL and Grid Engine (UGE) at T1_KIT
 - ◆ No meeting next week due to pre-GDB on batch systems
 - ◆ Then follow with reviews of torque/maui and SLURM
- Conclusions so far:
 - ◆ systems reviewed are capable of supporting multicore jobs

- ◆ however a tuning of each system is required to be able to absorb them (draining/reservation of resources) when running together with single core jobs
- ◆ a (so far) small degradation of CPU usage is noticed as a consequence of draining
- ◆ job submission pattern affects tuning, performance and wastage of the system. For ATLAS jobs:
 - ◇ pilots only running a single payload means that multicore slots don't survive long, therefore draining is constantly needed
 - ◇ wavelike pattern for multicore jobs creates the need to constantly tune the amount of draining needed
- Combined accounting of allocated and used resources for both single core and multicore jobs not clear so far
- Reminder that proper accounting requires APEL upgrade to EMI-3
- Discussion about multicore scheduling
 - ◆ The fact that batch systems examined so far release resources back to Condor pool and require renegotiation of the slot is a problem. The impact of draining depends on site size.
 - ◆ Simone comments that PanDA pilot framework does not support multiple payloads in one pilot, so not an option for ATLAS
 - ◆ Concerning job length, Alessandro comments that as short term mitigation "timefloor" can be increased in PanDA for multicore jobs.
 - ◆ Thomas comments that the site sees jobs arriving in bursts at intervals longer than job length. Alessandro and Simone comment that work is ongoing to fix 'burstiness' of multicore submission in PanDA
 - ◆ Alessandra and Antonio comment that batch systems could try to backfill, but requires experiment frameworks to provide wallclock information.

perfSONAR deployment TF

- Simone presents about perfSONAR deployment
 - ◆ <https://indico.cern.ch/event/280062/contribution/1/material/slides/?subContId=6>
 - ◆ perfSONAR 3.3.2 is now baseline
 - ◆ Task Force deadline is April 1st, all sites should have perfSONAR deployed, configured, registered, and opening ports in firewall for monitoring.
 - ◆ Prototype MaDDash monitoring available
- A few sites have valid concerns about opening firewall ports and require more restricted list of IP addresses, however it does not explain the large number of inaccessible sites.
- Alessandra and Simone remind that perfSONAR should be deployed "as close as possible" to storage, including same firewall configurations

SHA-2 Migration TF

- VOMRS
 - ◆ VOMRS was found to have become **compatible with SHA-2** when the VOMS clusters were upgraded to EMI-3 on Nov 27!
 - ◇ Many new users already registered OK with SHA-2 certificates.
 - ◆ Progress with the VOMS-Admin test cluster will now be tracked separately.
 - ◇ See the action list at the end of this page.
 - ◆ Host certs of our future VOMS servers are from the new SHA-2 CERN CA.
 - ◇ All VOMS-aware services on WLCG need to recognize the new servers before we can start using them.
 - ◇ We have prepared a campaign to be launched in the near future (not before next week).

- Maite comments that IT-PES wants to proceed anyway with VOMS-Admin commissioning since VOMRS is no longer supported. Progress is reported in the twiki linked in the action items.

Tracking tools evolution TF

- NTR

WMS decommissioning TF

- NTR

xrootd deployment TF

- NTR

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>

AOB

- The forum agrees to schedule the next Planning meeting on April 3rd.

This topic: LCG > WLCGOpsMinutes140306

Topic revision: r30 - 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