

Table of Contents

WLCG Operations Coordination Minutes, January 26th 2017.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Middleware News.....	1
Tier 0 News.....	2
Tier 1 Feedback.....	3
Tier 2 Feedback.....	3
Experiments Reports.....	3
ALICE.....	3
ATLAS.....	3
CMS.....	3
LHCb.....	4
Discussion.....	4
Ongoing Task Forces and Working Groups.....	4
Accounting TF.....	4
Information System Evolution.....	4
IPv6 Validation and Deployment TF.....	4
Machine/Job Features TF.....	4
Monitoring.....	5
MW Readiness WG.....	5
Network and Transfer Metrics WG.....	5
Squid Monitoring and HTTP Proxy Discovery TFs.....	5
Traceability and Isolation WG.....	5
Theme: Downtimes proposal followup.....	6
Theme: Tape usage performance analysis.....	6
Action list.....	6
Specific actions for experiments.....	7
Specific actions for sites.....	7
AOB.....	7

WLCG Operations Coordination Minutes, January 26th 2017

Highlights

- The new WLCG accounting portal [↗](#) has been validated.
- Please check the new accounting reports: if no major problems are reported, they become official as of Jan.
- Please check the baseline news and issues in the MW news.
- Long downtimes proposal and discussion.
- Tape staging test presentation and discussion.

Agenda

- <https://indico.cern.ch/event/607744/> [↗](#)

Attendance

- local: Alberto (monitoring), Alejandro (FTS), Alessandro (ATLAS), Andrea M (MW Officer + data management), Andrea S (IPv6), Andrew (LHCb + Manchester), Jérôme (T0), Julia (WLCG), Kate (WLCG + databases), Maarten (WLCG + ALICE), Maria (FTS), Marian (networks + SAM), Vincent (security)
- remote: Alessandra (ATLAS + Manchester), Catherine (IN2P3 + LPSC), Christoph (CMS), David B (IN2P3-CC), Di (TRIUMF), Frédérique (IN2P3 + LAPP), Gareth (RAL), Kyle (OSG), Marcelo (LHCb), Oliver (data management), Renaud (IN2P3-CC), Ron (NLT1), Stephan (CMS), Thomas (DESY-HH), Vincenzo (EGI), Xin (BNL)
- Apologies: Nurcan (ATLAS), Ulf (NDGF-T1)

Operations News

- Next WLCG Ops Coord meeting will be on March 2nd

Middleware News

- Useful Links:
 - ◆ Baseline Versions [↗](#)
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines/News:
 - ◆ New APEL 1.4.1 and APEL-SSM 2.17 supporting SL7/C7 (GGUS:126009 [↗](#)). To be included in UMD 4
 - ◆ new ARC-CE release (http://www.nordugrid.org/arc/releases/15.03u11/release_notes_15.03u11.html [↗](#)) available also in EPEL. It fixes an issue with HTCondor 8.5.5 reported by some sites affecting the job status query.
 - ◆ dCache 2.10.x support ended in 2016. From the BDII we still see 18 instances running this version (including BNL) . We will coordinate with EGI in order to ticket the sites asking to upgrade to 2.16
- Issues:

- ◆ Follow up on the issue reported by RFC proxy TF (GGUS:124650). A new version of dCache 2.13 has been released last week (2.13.51, <https://www.dcache.org/downloads/1.9/release-notes-2.13.shtml>). It fixes a problem in jglobus discovered in the context of the RFC Proxy TF, which prevented RFC proxies of certificates belonging to certain CAs to work with dCache 2.13 (from v 2.14 jglobus has been removed from dCache). The same issue is still present in Bestman though.
 - ◆ SLURM high risk vulnerability (<https://wiki.ege.eu/wiki/SVG:Advisory-SVG-CVE-2016-10030>). Sites are requested to update to version 15.08.13 or 16.05/8
 - ◆ High risk CVE-2016-7117 Linux kernel vulnerability (<https://wiki.ege.eu/wiki/SVG:Advisory-SVG-CVE-2016-7117>). Patched kernels are available for RHEL 6/7 and SL6/7
- T0 and T1 services
 - ◆ CERN
 - ◇ check T0 report
 - ◇ FTS upgrade to v. 3.5.8 and C7.3 planned for next week
 - ◆ FNAL
 - ◇ FTS upgrade to v. 3.5.7
 - ◆ IN2P3
 - ◇ dCache upgrade 2.13.32 -> 2.13.49 on Dec 2016
 - ◆ JINR
 - ◇ dCache minor upgrade 2.13.49 -> 2.13.51, xrootd minor upgrade 4.4.0 -> 4.5.0-2.osg33
 - ◆ NDGF-T1
 - ◇ Upgrade to dCache (3.0.5) next week to fix a communication bug that is rare.
 - ◆ NL-T1
 - ◇ SURFsara upgraded dCache from 2.13.29 to 2.13.49 on Dec 1
 - ◆ RAL
 - ◇ Castor upgrade to v 2.1.15-20 ongoing, Tape servers upgraded to v 2.1.16
 - ◇ Almost completed migration of LHCb data to T10KD drives.
 - ◇ Update SRMs to version 2.1.16-10 has been planned
 - ◆ RRC-KI-T1
 - ◇ Upgraded dCache for tape instance to v 2.16.18-1

Tier 0 News

- Apex 5.1, released in December, has been installed on the development instance; the upgrade of the production environments is scheduled for February.
- The main compute services for the Tier-0 are now provided by 31k cores in HTCondor, 70k cores in the LSF main instance, 12k cores in the dedicated ATLAS Tier-0 instance, and 21.5k cores in the CMST0 compute cloud. Some CC7-based capacity is available in HTCondor for user testing. We are in contact with major user groups for consultancy for migrating to HTCondor.
- 2016 LHC data taking ended with the p-Pb run; about 5 PB have been collected in December. Since then there has been a lot of consolidation activity.
- The p-Pb run for ALICE was performed using a 1.5 PB Ceph-based staging area, without incident or slow-down. This configuration, offering increased flexibility and easier operation, is now considered production-ready for CASTOR.
- Mostly transparently to users, EOS and CASTOR instances have been rebooted. On 23 January the failover of an EOS ATLAS headnode failed, causing service disruption; the root cause is being investigated. Some EOS CMS instability requiring a headnode to be rebooted has been due to heavy user activity.
- New storage hardware will be installed in February as soon as it becomes available to the service.

- The FTS service has been upgraded to v.3.5.8 and a refreshed VM image based on CentOS 7.3 fixing vulnerabilities.
- In January the mark of one billion files in EOS at CERN was crossed.

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- High to very high activity since the end of the proton-ion run
 - ◆ Also during the end-of-year break
 - ◆ In particular for Quark Matter 2017 [☞](#), Feb 5-11
 - ◆ Thanks to the sites for their performance and support!
- On Wed Jan 4 the AliEn central services suffered a scheduled power cut
 - ◆ Normally such cuts are short and do not present a problem
 - ◆ This time the UPS was exhausted many hours before the power was restored
 - ◆ All ALICE grid activity started draining away
 - ◆ Fall-out from the incident took ~2 days to resolve
- A recent MC production is using very high amounts of memory
 - ◆ Large numbers of jobs ended up killed by the batch system at various sites
 - ◆ Experts are looking into reducing the memory footprint further

ATLAS

- Very smooth production operations during the winter break; MC simulation, derivations and user analysis have been running in full speed since then. Derivations use up to 100k slots to finish full processing of data15+data16 and their MC samples. Very high user analysis activity for Moriond and spring conferences.
- ATLAS Sites Jamboree took place at CERN on January 18-20. Good feedback and discussions from sites. Interest on virtualization from sites; docker containers and singularity, a dedicated meeting is being planned before Software&Computing week in March.
- Global shares are being implemented to better manage the required resources among different workflows (MC, derivations, reprocessing, HLT processing, Upgrade processing, analysis). More intensive productions will run in the future. Sites will be asked to get rid of the fairshares set at the site level. This is currently in revision and some sites that ran only evgensimul (or part of it) in the past will be affected.
- Tape staging test at 3 Tier1 s two weeks ago to understand the performance in case running the derivations from the tape input might become a possibility. Results will be presented today:
20170126_Tape_Staging_Test.pdf

CMS

- good, steady progress on Monte Carlo production for Moriond 2017
 - ◆ tape backlog at sites worked down
 - ◆ JINR looking into tape setup to improve performance
- PhEDEx agents were updated in time at all but two sites, ready to switch to FTS3 and new API
 - ◆ *Christoph, Stephan: the old SOAP interface can be switched off now*
 - ◆ *FTS team: will be done in the intervention next Tue*
- dedicated hi-performance VMs for GlideinWMS, under evaluation using Global Pool scalability test

- slots overfilled at sites due to HT Condor bug (triggered by node with high network utilization/errors); bug will be addressed in v8.6 to be released early February
 - ◆ *Stephan: a fraction of the jobs used more cores than reserved; so far this could be mitigated*
- CentOS 7 plans: no general upgrade planned; sites are free to upgrade and provide SL6 via container, etc.; CMS software is also released for CentOS 7; physics validation expected soon; discussions and tests to move HLT farm to CentOS 7;
- pilots sent to Tier-1 sites consolidated: one pilot with role "pilot" instead of two with roles "pilot" and "production"

LHCb

- *Andrew: high activity, nothing special to report*

Discussion

- David B: what is the general readiness for CentOS7 WNs?
- Maarten: ALICE and LHCb can use such resources today
- Christoph, Stephan: the CMS report has the details for CMS
- Alessandro: ATLAS production can run OK, analysis to be checked; we will check if SL6 containers are OK
- *Alessandra after the meeting: the situation for ATLAS was presented [in the ATLAS Sites Jamboree](#) held on Jan 18-20*
 - ◆ *All SL6 SW workflows have been validated*
 - ◆ *Site admins should read the presentation for details, considerations and advice*

Ongoing Task Forces and Working Groups

Accounting TF

- New WLCG accounting portal (<http://accounting-next.egi.eu/wlcg>) has been validated and people are welcome to start using it
- Migration to the new accounting reports started. Two sets of accounting reports (current and new ones) have been sent for November and December. In case no major problems are reported, starting from January, new reports become official.

Information System Evolution

- Next meeting Feb 2.

IPv6 Validation and Deployment TF

- NTR

Machine/Job Features TF

- Further updates to DB12-at-boot in mjf-scripts distribution. See discussion at HEPiX benchmarking WG meeting. These values are made available as \$MACHINEFEATURES/db12 and \$JOBFEATURES/db12_job

Monitoring

- NTR

MW Readiness WG

This is the status of Actions from the 19th MW Readiness WG meeting of 20161102.

- 20161102-05: Christoph to investigate EL7 UI testing by CMS. Keep Andrea S. informed as maintainer of the workflow twiki.
- 20161102-01: Andrea S. to update the CMS workflow twiki.

This is the status of jira ticket updates [↗](#) since the last Ops Coord of 20161201:

- MWREADY-138 [↗](#) DPM 1.9.0 on C7 at GRIF_LL2 for CMS - completed
- MWREADY-104 [↗](#) DPM 1.9.0 SRM-less for ATLAS at LAPP Annecy - on-going. Atlas verified the modification to the pilot to use dav for stage-in/out
- MWREADY-142 [↗](#) FTS 3.5.8 for ATLAS & CMS at CERN - on-going
- MWREADY-140 [↗](#) ARC-CE 5.2.1 on C7 for CMS at Brunel - on-going (ARC-CE 5.2.0 completed)
- MWREADY-135 [↗](#) WN for C7/SL7 at TRIUMF for ATLAS - on-going
- MWREADY-141 [↗](#) dCache 3.0.2 at PIC for CMS - on-going

Network and Transfer Metrics WG

- pre-GDB on NETWORKING took place on 10th of January (<https://indico.cern.ch/event/571501/>)
 - ◆ Summary was presented at the GDB next day (<https://indico.cern.ch/event/578982/>)
- Throughput meeting held on 14th of Dec (<https://indico.cern.ch/event/595286/>)
 - ◆ Discussed ongoing validation of perfSONAR 4.0; RC3 to be released next week
- perfSONAR developers meeting ongoing this week
- WLCG Network Throughput Support Unit: see twiki for summary of recent activities.

Squid Monitoring and HTTP Proxy Discovery TFs

- CERN now has the first site-specific <http://grid-wpad/wpad.dat> [↗](#) service, and CMS is using it for their jobs
 - ◆ hosted on the same 4 physical 10gbit/s servers that CMS uses for squid service
 - ◆ supports both IPv4 and IPv6 connections in order to determine whether squids in Wigner or Meyrin should be used first
 - ◆ for CMS destinations (Frontier or cvmfs), it directs to the CMS squids, otherwise it defaults to the IT squids

Traceability and Isolation WG

- A tool was identified as a possible solution for isolation (without traceability), 'singularity':
 - ◆ WG now evaluating the tool: a small test cluster is being built now
 - ◆ A security review of this tool is needed (transition plan might require SUID)
- Next meeting: Wednesday 1 Feb 2017 (<https://indico.cern.ch/event/604836/>)

Theme: Downtimes proposal followup

See the presentation [↗](#)

- Andrew: shifting a downtime can upset the experiment planning, it should not be done lightly
- Stephan: only explicit policies can be programmed automatically
- Maarten: we can have SAM apply the numbers of the new policies, whereas the clauses in purple (see presentation) would be "best practice"; we can always do manual corrections as needed
- Julia: manual operations should only be done in exceptional cases, e.g. when a downtime had to be extended or postponed
- Maarten: today's feedback plus the outcome of a discussion with the SAM team will be incorporated into v3 to be presented in the MB

Theme: Tape usage performance analysis

See the presentation

- Alessandro: can the FTS handle batches of 90k requests per T1, times 5 sites?
- FTS team: v3.6 should be able to handle 500k per link
- CMS: we want to understand the tape usage performance through a similar exercise; afterwards we can do a common challenge together with ATLAS
- Alessandro: the common challenge would be at a shared T1
- CMS: the monitoring is tricky; a file could first be read from tape and then multiple times from disk; some files could already be on disk; such cases would have to be disentangled
- Alessandro: that is why in the ATLAS exercise we took very old files
- Alessandro: the throughput was measured between submission of the request and the file being present in the ATLAS_DATA_DISK space
- Alessandro: new files could be made bigger (not trivial), existing files cannot
- Renaud: w.r.t. the number of files per request, we will check what limitations exist on our side
- Alessandro: a typical campaign would cover O(100) TB and O(100 k) files; we can customize that per site
- Renaud: what do you think of the observed performance?
- Alessandro: it did not change over the past years, because experiments did not work on it ! it also depends on site-specific tape handling aspects
- Julia: shouldn't also the recording be optimized?
- Alessandro: indeed, e.g. through new tape families

Action list

Creation date	Description	Responsible	Status	Comments
01 Sep 2016	Collect plans from sites to move to EL7	WLCG Operations	Ongoing	The EL7 WN is ready (see MW report of 29.09.2016). ALICE and LHCb can use it. NDGF plan to use EL7 for new HW as of early 2017. Other ATLAS sites e.g. Triumf are working on a container solution that could mask the EL7 env. for the experiments which can't use it. Maria said that GGUS tickets are a clear way to collect the sites' intentions. Alessandra said we shouldn't ask a vague question. Andrea M. said the UI bundle is also making progress.

				Jan 26 update: this matter is tied to the EL7 validation statuses for ATLAS and CMS, which are reported above.
03 Nov 2016	Review VO ID Card documentation and make sure it is suitable for multicore	WLCG Operations	Pending	Jan 26 update: needs to be done in collaboration with EGI
03 Nov 2016	Discuss internally on how to follow up long term strategy on experiments data management as raised by ATLAS	WLCG Operations	DONE	Jan 26 update: merged with the next action
03 Nov 2016	Check status, action items and reporting channels of the Data Management Working Group	WLCG Operations	Pending	
26 Jan 2017	Create long-downtimes proposal v3 and present it to the MB	WLCG Operations	Pending	

Specific actions for experiments

Creation date	Description	Affected VO	Affected TF/WG	Comments	Deadline	Completion
29 Apr 2016	Unify HTCondor CE type name in experiments VOfeeds	all	InfoSys	Proposal to use HTCONDOR-CE.		Ongoing
01 Dec 2016	Open tickets to sites for moving to FTS3 client	CMS	-	There are PhEDEx prerequisites	Jan 2017	DONE

Specific actions for sites

Creation date	Description	Affected VO	Affected TF/WG	Comments	Deadline	Completion
01 Dec 2016	Proposal for advance warning of long site downtimes	All	-	Please, give feedback to the Dec GDB proposal ↗	20th January 2017	DONE

AOB

This topic: LCG > WLCGOpsMinutes170126

Topic revision: r20 - 2018-02-28 - MaartenLitmaath



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback