

Table of Contents

WLCG Operations Coordination Minutes, Dec 7th 2017.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	2
Middleware News.....	2
Discussion.....	3
Tier 0 News.....	3
Tier 1 Feedback.....	3
Tier 2 Feedback.....	3
Experiments Reports.....	3
ALICE.....	3
ATLAS.....	4
CMS.....	4
LHCb.....	5
Ongoing Task Forces and Working Groups.....	5
Accounting TF.....	5
Archival Storage WG.....	5
Information System Evolution TF.....	5
IPv6 Validation and Deployment TF.....	5
Machine/Job Features TF.....	6
Monitoring.....	6
MW Readiness WG.....	6
Network Throughput WG.....	6
Squid Monitoring and HTTP Proxy Discovery TFs.....	6
Traceability WG.....	7
Container WG.....	7
Special topics.....	7
Follow up on CNAF incident.....	7
Update from Container WG.....	7
Update from Archival Storage WG.....	8
SAM and SSB migration to MONIT.....	8
Action list.....	9
Specific actions for experiments.....	9
Specific actions for sites.....	9
AOB.....	9

WLCG Operations Coordination Minutes, Dec 7th 2017

Highlights

- The CMS experiment plans to process data from the upcoming LHC run in 2018 using SL/CentOS 7. Analysis of existing data will continue to use SL 6. The WLCG Container working group is looking into Singularity for job OS flexibility and for job isolation. Singularity can be installed on either SL 6 or CentOS 7 worker nodes from the OSG or WLCG repository. CMS recommends CentOS 7 plus Singularity. Documentation with a link to OSG setup instructions can be found [here](#). With Singularity glExec is also no longer needed/used by CMS. CMS plans to include the SAM Singularity probe in the critical profile in March. SAM site availability will then drop to zero if Singularity is not installed !
 - ◆ WLCG operations will broadcast further information to the sites.
 - ◆ Also see the update from the Container WG linked below.
- More details regarding follow up on CNAF incident
 - ◆ To summarize the discussion regarding followup on CNAF incident we can say that the impact was less serious than it could have been thanks to the redundancy of the LHC computing models. No data loss for data which cannot be reproduced. None of the experiments stated that there was a disaster. Lack of CPU resources was easily compensated by other sites. Only CMS mentioned that they asked other T1 sites to increase their shares. None of the experiments mentioned that users had been heavily affected and complained. The most painful problem is to reproduce samples which were lost. If experiments could avoid having a single copy of reconstructed data, imposed by RRB constraints, the recovery would be much easier. It would be useful to estimate how expensive the recovery was regarding involved manpower. This can be estimated later, most probably in February.
- Update from Container WG
- Update from Archival Storage WG
- SAM and SSB migration to MONIT

Agenda

- <https://indico.cern.ch/event/684980/>

Attendance

- local: Alberto (monitoring), Alessandro Di G (ATLAS), Andrea M (MW Officer + data management), Antonio (CNAF + LHCb), Christoph (CMS), Gavin (T0), Giuseppe B (CMS), Giuseppe Lo P (storage), Ivan (ATLAS), Julia (WLCG), Kate (DB + WLCG), Maarten (ALICE + WLCG), Marian (networks), Oliver (data management), Stefan (LHCb)
- remote: Alessandra (Manchester + ATLAS), Alessandro (?), Andrea S (WLCG), Catherine (LPSC + IN2P3), Daniele (CNAF), Dave D (FNAL), Di (TRIUMF), Eddie (monitoring), Edith (Marseille), Eric F (IN2P3-CC), Eric L (BNL), Felix (ASGC), Fred (GRIF-IRFU), Kyle (OSG), Luca (monitoring), Mayank (WLCG), Ron (NLT1), Stephan (CMS), Ulf (NDGF), Victor (GRIF-LPNHE), Vincenzo (EGI), Xin (BNL)

- apologies:

Operations News

- January meeting will be dedicated to site availability reports and SAM tests. We would like to assess usefulness of these reports (for sites, VOs, WLCG managements) and whether critical profiles do evaluate properly site availability and whether they are used for operations do detect problems with the sites. We will require input from the experiments and sites.
- The next meeting is planned for Jan **18**
 - ◆ Please let us know if that date would present a significant problem

Middleware News

- Useful Links:
 - ◆ Baseline Versions [↗](#)
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines/News
 - ◆ As requested by LHCb, new versions of HEP_OSlibs for EL7 (7.2.3-1) and HEP_OSlibs_SL6 (1.1.2.-1) are available on the WLCG repos. Site supporting LHCb VO should upgrade their WNs.
- Issues:
 - ◆ Issue with several services after new UK Science CA has been released in lcg-CA 1.88.1. It seems that the change to SHA256 of that CA triggered an issue in CANL-Java which is used by dCache, so clients having that CA certs are not able to contact dCache SRM and gridftp doors which have upgraded to lcg-CA 1.88.1. (see GGUS:132205 [↗](#), GGUS:132222 [↗](#), GGUS:132237 [↗](#)) Paul Millar from dCache has a proposed fix for CANL-Java already (<https://github.com/eu-emi/canl-java/pull/93> [↗](#)), but for now there is a workaround which is to restart all dCache GSI based doors.
 - ◆ The same issue is affecting other JAVA based components, like VOMS-Admin. Still further investigation is needed to understand the impact.
 - ◆ Several DPM SRM crashes reported last week by ATLAS sites. They all happened after being contacted by a specific user DN and IP. We identified the issue and a new DPM release is under preparation.
 - ◆ Vulnerability broadcasted by SVG. (<https://wiki.eui.eu/wiki/SVG:Advisory-SVG-CVE-2017-5712> [↗](#)) Various vulnerabilities have been identified by Intel, which allow privilege escalation. At least one site has been identified which has the vulnerable CPU type in the cluster
- T0 and T1 services
 - ◆ BNL
 - ◆ FTS upgrade to 3.7.7
 - ◆ CERN
 - ◆ check T0 report
 - ◆ EOS Alice and LHCb upgraded to 4.2.2
 - ◆ FTS upgraded to v 3.7.7
 - ◆ IN2P3
 - ◆ Update of dcache from 2.16.46 to 2.16.54
 - ◆ JINR
 - ◆ Minor upgrade: dCache 2.16.51->2.16.53; Postgres 9.6.5->9.6.6; upgrade xrootd 4.6.1->4.7.1
 - ◆ KIT:
 - ◆ Updated dCache to 2.16.54 for LHCb (Nov 28th), ATLAS (Nov 29th) and CMS

(Nov 29th).

- ◆ RAL:
 - ◇ FTS upgrade to 3.7.7 (both instances)

Discussion

- Maarten: we will follow up offline to come up with a broadcast about the UK CA issue
- Stephan: it would be good to broadcast preliminary info earlier next time
- Alessandra: the UK community was badly affected
 - ◆ *Added after the meeting:*
 - ◇ https://www.gridpp.ac.uk/wiki/Problems_After_CA_1.88-1
- Maarten:
 - ◆ we were waiting for contradictory statements to get sorted out
 - ◆ nobody asked for a broadcast
 - ◆ next time we could broadcast early that the matter is being investigated
- Andrea M:
 - ◆ next time please involve WLCG Operations earlier
 - ◆ this time we learned about the problems a bit accidentally...
- Alessandra: another issue is that sites could not easily roll back, because IGTF do not keep the previous rpms available
- Maarten: we will follow up with them
- *Added after the meeting:*
 - ◆ *The old rpms actually are available in at least 2 places:*
 - ◇ <https://egi-igtf.ndpf.info/distribution/>
 - ◇ <https://dist.igtf.net/distribution/igtf/>
 - ◆ *However, it still looks not completely trivial to roll back...*
- Andrea, Maarten: ideally we would have MW readiness validation also for CAs, but WLCG does not have the resources for that
 - ◆ As such problems are rare, we can just take the hit...

Tier 0 News

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- The activity levels have been high to very high on average.
 - ◆ Up to 162k concurrent jobs overall, 84k at CERN.
- CERN:
 - ◆ CASTOR has been enlarged with 2 PB borrowed from EOS-ALICE.
 - ◇ That allowed the data taking to finish successfully.
 - ◆ CASTOR then was reconfigured to allow large-scale reco.
 - ◇ Up to 50k concurrent jobs have been running OK.
 - ◆ EOS: severe instabilities and downtime on Nov 10 (GGUS:131749)
 - ◇ ~1.2 M files were temporarily lost from the name space.
 - ◆ Both CASTOR and EOS saw a few other instabilities as well.

- ◆ We thank the storage group for their much appreciated efforts!

- **CNAF flooding - effects on ALICE**

- ◆ 62 M files comprising 3 PB unavailable from the disk SE.
- ◆ The most important ones have replicas elsewhere.
- ◆ A few hundred TB are single-replica ESD files.
- ◆ Other files (notably AODs) are being re-replicated between other SEs.
- ◆ To replicate all would take a few months.
- ◆ There are 20 wet tapes with ALICE data: they do **not** need to be recovered.
 - ◇ Their data will be re-replicated from tapes at CERN.

ATLAS

- Stable grid production operations with ~300k concurrent running job slots. Good HPC contributions compared to the last month with peaks up to 600k jobs.
- Derivation campaign (as a reminder, high I/O jobs, inputs order of 3-4PB, output 8-9PB) started 2 weeks ago, running smoothly at ~100k slots in parallel (peaks at 120), forecast it will last another 3-4 weeks
- Reprocessing campaign ready to be launched (this Friday). Expected duration is about 3-4 weeks (for the first half of 2017 data, the second half will be done in January). Inputs are on tape, we won't pre-stage them in advance (as we did e.g. for the reprocessing we run in April-July), we will see how the automatic prestage will work.
- Last runs of 2017 with very high rates (more than continuous 2 GB/s, one run in particular lasted 36hours and was recording 3GB/s data) partially hit limits of EOS/Castor export rate which are still under investigation by ATLAS DDM and CERN IT experts. It seems that EOS GridFTP servers (9) are kind of saturated. But also it's not clear why in general we are not able to push more into CASTOR (we saw peaks up to 3GB/s, maybe we are not able to keep the queue always full?).
- Sim@P1 (ATLAS high level trigger farm resources) are available for Simulation, providing ~45k slots more.

CMS

- Tier-0 reached its limit
 - ◆ transfer backlog dropped below 5 PB now that data taking stopped
 - ◆ CPU resources used now by production
- Compute systems are busy
 - ◆ 2017 data re-reco started
 - ◆ Monte Carlo samples for 2017 analyses
 - ◆ Phase-II upgrade simulations
- plan to run 2018 data processing on SL/CentOS 7
 - ◆ CMS would like to use Singularity at sites end of 2017 or by March of 2018 the latest (so we can run both SL 6 and 7 in parallel)
 - ◆ SL 7 with Singularity recommended (SL 6 + Singularity also ok)
 - ◆ what is WLCG's Singularity position?
 - ◆ can/will WLCG coordinate with grid organizations/sites the Singularity roll-out?
 - ◆ SAM probe for Singularity already running, test will then become mandatory, i.e. enter site availability as of March
- recent production workflows exceeded address space/virtual memory requirement in VO-card (due to jmalloc use, no change in actual memory use just address space squandering); VO-card will be updated to 8 GB/core
- CMS report about CNAF incident [link](#)

LHCb

- Stripping campaign ongoing stressing tape system to recall Reco15 Reco16 data
- On CNAF incident:
 - ◆ CPU is not a main issue, mitigated by the help of other sites
 - ◆ 27 Tape affected recoverable from CERN copy, RDST/FULL.DST single copy to be addressed after CNAF will be online
 - ◆ Disk will delay Run 2 stripping campaign

Ongoing Task Forces and Working Groups

Accounting TF

- EOS, DPM, STORM and dCache development teams confirmed that they were progressing in implementation of the Storage Resource Reporting Proposal which has been approved by November GDB and which is required for the WLCG Storage Accounting implementation. XRootD also confirmed that they would do it.

Archival Storage WG

- First meeting held - <https://indico.cern.ch/event/681096/>
 - ◆ Minutes attached to agenda
- Twiki space revived
 - ◆ <https://twiki.cern.ch/twiki/bin/view/HEPTape/WebHome>
- First draft of questionnaire
 - ◆ <https://twiki.cern.ch/twiki/bin/view/HEPTape/Recommendations>
 - ◆ Aim is to compile "best practice" on exploitation of tape resources
- Agreement to proceed with metric sharing
 - ◆ Vlado to identify platform and first metrics

Information System Evolution TF

- At the November IS Evolution Task Force meeting [discussed](#) static/semi-static and dynamic info required for compute resource description. The discussion was initiated by the fact that sites running ARC had to apply custom patches in order to provide dynamic info in BDII. Dynamic information, like number of pending and running jobs is currently used by ALICE and LHCb. They both confirmed that they are fine with GLUE2. The decision was to request the central fix of the ARC problem by ARC development team. Apparently, if the fix is being limited to GLUE2 only, it should be quite straight forward. The ticket for this problem has been submitted a while ago. Before the fix is provided, sites should not develop and apply any custom patches. As a common strategy, both LHCb and ALICE will use dynamic info from BDII, unless the site does not run BDII. In case when BDII is lacking, both LHCb and ALICE confirmed that they have a way around.
- Description of the compute resources in CRIC has been implemented. Started implementation of the storage resource description.

IPv6 Validation and Deployment TF

Almost all Tier-2's have received a GGUS ticket to keep track of their IPv6 deployment status. More info [here](#).

- Andrea S:
 - ◆ 11% of the T2 sites have IPv6 deployed and tested
 - ◆ many sites expect to deploy it in the next months

- ◆ so far no site has said it cannot meet the deadline
- ◆ by the end of next year the situation is expected to look quite good

Machine/Job Features TF

Monitoring

MW Readiness WG

NTR

Network Throughput WG

- perfSONAR workshop held by JISC in November (slides [↗](#)) - WLCG WG activities mentioned
- perfSONAR 4.0.2 was released on November 28th
 - ◆ **WLCG broadcast** will be sent this week to notify sites of the **upcoming important dates** and new documentation
 - ◆ perfSONAR 4.1 release, planned in Q1 2018 will no longer ship SL6 packages
 - ◆ EOL for SL6 support in Q3 2018
 - ◆ All sites are encouraged to upgrade to CC7 as soon as possible
- WLCG/OSG network services
 - ◆ On-going work on improving OSG collector and broadcast results via GOC's RabbitMQ
 - ◆ New documentation is available at <https://opensciencegrid.github.io/networking/> [↗](#)
 - ◆ New LHCPN and perfSONAR dashboards done in collaboration with OSG, UC, CERN IT/CS and IT/MONIT are available at <http://monit-grafana-open.cern.ch/dashboard/db/home?orgId=16> [↗](#)
- WLCG Network Throughput Support Unit: see twiki for summary of recent activities.
- HNSciCloud will use perfSONAR results for network performance evaluation of the providers
- **HEPiX WG on SDN/NFV** was established and will look into networking R&D topics - sites interested please subscribe via <https://listserv.in2p3.fr/cgi-bin/wa?SUBED1=hepixonfv-wg> [↗](#)
 - ◆ WG scope and objectives yet to be defined, Doodle for an initial meeting will be sent soon
 - ◆ WG proposal presentation is at <https://indico.cern.ch/event/637013/contributions/2739266/> [↗](#)

Squid Monitoring and HTTP Proxy Discovery TFs

- All ATLAS Tier 1 and Tier 2 sites are now having their squids monitored with MRTG, almost all via being registered in GOCDB or OIM plus a small number added by exception. ATLAS now has its own MRTG monitoring page [↗](#) with an improved user interface.
- CMS OSG Opportunistic jobs are now using <http://openhtc.io> [↗](#) aliases for CMS Frontier servers. Those send all the traffic through the free tier of the Cloudflare Content Delivery Network, which has unlimited bandwidth.
 - ◆ Web Proxy Auto Discovery has been updated so that if the server is in openhtc.io, and no WLCG squids found, it will use DIRECT to the server
 - ◆ In order to make that work, had to remove backup proxies from the configuration. WLCG sites are **not** set to failover to DIRECT because then the squids could be failing and nobody would notice, so it was at first set up to instead cause jobs to fail if the squids failed.
 - ◆ There has not yet been any attempted usage at sites without WLCG squids, but the Syracuse University squid, which is in the WLCG list via an exception, has been having problems this past week and over the weekend it was removed from the WLCG squid list and all worker nodes directed to Cloudflare
 - ◇ Held up well, handled about 750K requests per hour

- ◆ Today WLCG WPAD started returning backup proxies at the end of every list of proxies, which is tricky because there are different backup proxies for different services
- We are working on converting LHC@Home to also use <http://openhtc.io> aliases for CMS Frontier servers and CVMFS stratum 1 servers.
- Dave D: we probably are below Cloudflare's "noise threshold" for the time being, but we have to see how long that will last

Traceability WG

Container WG

See update slides.

Special topics

Follow up on CNAF incident.

See the presentations [↗](#)

- Julia: any lessons for LHCb besides the observed benefits from tape families?
- Stefan: it would be best to avoid single copies for reco output, but we were forced to do that as RRB C-RSG discussions concluded that it was a good way to save on disk space consumption...
- Antonio: what data sets will be reproduced?
- Ivan: only urgent ones requested by the physicists
- Julia: were there many user complaints about unavailable data?
- Giuseppe B: not AFAIK
- Julia: were compensations by other T1 + T0 also offered to other experiments?
- Various: yes
- Julia: what was the main lesson for CMS?
- Giuseppe B: try to avoid single copies to the extent possible, but there is the same RRB counter argument as explained by Stefan
- Maarten: ditto for ALICE
- Maarten: the impact on ALICE operations seems to be under control so far, but users did report issues due to the unavailability of files at CNAF
- Alessandro Di G: ATLAS has tried to quantify how expensive the recovery has been in terms of FTEs; it would be good if the other experiments could do the same, so that together we may present to the RRB C-RSG the associated costs in comparison with the cost of extra disk space
- Stefan: LHCb will only know the full impact in Feb or March
- Julia: we may follow up through the GDB etc.

Update from Container WG

See the presentation [↗](#)

- Alessandra: we should expand on the CMS request to have Singularity deployed at all their sites by March.
- Christoph:
 - ◆ the 2018 data will be processed only with EL7 releases of CMSSW
 - ◆ for older data we still need SL6 environments
 - ◆ sites may upgrade to EL7 and provide Singularity for SL6 jobs
 - ◆ or they can stay at SL6 and provide Singularity for EL7 jobs
- Julia: what would be the deployment plan?
- Giuseppe B: we have informed our sites, but a message from WLCG would be good
- Julia: do you check the sites?
- Christoph: we have a SAM test for that and we have meetings involving sites
- Maarten:
 - ◆ let's follow up offline to arrive at a WLCG broadcast about this
 - ◆ we need a pointer to instructions from CMS
- Stephan:
 - ◆ in March we will make that SAM test critical and open tickets as needed
 - ◆ until that time the test just flags if sites look OK or not
 - ◆ we also need to understand if there is any conflict with WLCG plans?
- Alessandra:
 - ◆ Making SAM tests critical needs to be pointed out in the broadcast
- Gavin:
 - ◆ we track different requirements from the experiments in the WG
 - ◆ we will try to come up with a single set if possible

Update from Archival Storage WG

See the update

No comments.

SAM and SSB migration to MONIT

See the presentation [↗](#)

- Julia: why is HDFS needed for certain queries?
- Alberto:
 - ◆ Elasticsearch only holds the last month
 - ◆ InfluxDB just has aggregates for the last 5 years
 - ◆ HDFS has everything forever
- Alessandra: can one select a range of days?
- Alberto: yes
- Alessandra: can one dig down to the errors?
- Alberto:
 - ◆ in Grafana that is easy for some cases, expensive for others
 - ◆ we will iterate with experts until the functionality looks sufficient
- Julia: as validation by eye is difficult, will you automate it?
- Alberto: we plan for that, but it will take time
- Maarten: the plots do not have to look exactly the same as in SAM-3, they just need to be usable

- Julia: will you implement derived metrics?
- Alberto: to be seen

- Maarten:
 - ◆ ALICE relies on VO-injected metrics for the critical profile
 - ◆ that feature is not a luxury we can do later

- Julia:
 - ◆ WLCG relies a lot on SSB functionality
 - ◆ we can volunteer some effort for evaluating and testing of new SSB implementation

Action list

Creation date	Description	Responsible	Status	Comments
01 Sep 2016	Collect plans from sites to move to EL7	WLCG Operations	Ongoing	[older comments suppressed] Sep 14 update: UMD 4.5 was released on Aug 10, containing the WN; CREAM and the UI are available in the UMD Preview repo; CREAM client tested OK by ALICE Dec 7 update: Tier-1 plans are documented in the Nov 2 minutes
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
26 Jan 2017	Create long-downtimes proposal v3 and present it to the MB	WLCG Operations	In progress	May 18 update: EGI collected feedback from sites and propose a compromise - 3 days' notice for any scheduled downtime. Oct 5 update: as both OSG and EGI were not happy with the previous proposals, and as this matter does not look critical, we propose to create best-practice recipes instead and advertise them on a WLCG page. Dec 7 update: a first version is prominently linked from WLCGOpsWeb
14 Sep 2017	Followup of CVMFS configuration changes, check effects on sites in Asia	WLCG Operations	Pending	

Specific actions for experiments

Creation date	Description	Affected VO	Affected TF/WG	Comments	Deadline	Completion
---------------	-------------	-------------	----------------	----------	----------	------------

Specific actions for sites

Creation date	Description	Affected VO	Affected TF/WG	Comments	Deadline	Completion
---------------	-------------	-------------	----------------	----------	----------	------------

AOB

** Season's greetings and best wishes for the New Year !*

-- JuliaAndreeva - 2017-12-01

This topic: LCG > WLCGOpsMinutes171207

Topic revision: r29 - 2017-12-13 - MaartenLitmaath



Copyright &© 2008-2020 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](#)