

Table of Contents

WLCG Operations Coordination Minutes, April 12th, 2018.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
SAM recalculation policy.....	1
Discussion.....	1
Middleware News.....	2
Discussion.....	2
Tier 0 News.....	2
Tier 1 Feedback.....	2
Tier 2 Feedback.....	2
Experiments Reports.....	2
ALICE.....	2
ATLAS.....	3
Discussion.....	3
CMS.....	3
LHCb.....	4
Discussion.....	4
Ongoing Task Forces and Working Groups.....	4
Accounting TF.....	4
Archival Storage WG.....	5
Update of providing tape info.....	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.....	6
Container WG.....	6
Special topics.....	6
Action list.....	6
Specific actions for experiments.....	7
Specific actions for sites.....	7
AOB.....	7

WLCG Operations Coordination Minutes, April 12th, 2018

Highlights

- SAM recalculation policy

Agenda

- <https://indico.cern.ch/event/711927/>

Attendance

- local: Alberto (monitoring), Alessandro (ATLAS), Borja (monitoring), Christoph (CMS), Dimitrios (WLCG), Federico (LHCb), Giuseppe (CMS), Julia (WLCG), Maarten (WLCG + ALICE), Mayank (WLCG), Stephan (CMS)
- remote: Alessandra D (Napoli), Alessandra F (Manchester + ATLAS), Balazs (MW Officer), Di (TRIUMF), Jeremy (GridPP), Ricardo (SAMPa), Ron (NLT1), Thomas (DESY), Xin (BNL)
- apologies:

Operations News

- T1 sites are requested to complete the tape survey
- We welcome Balazs Konya in a new role of WLCG Middleware Officer!
- The next meeting is planned for **June 7**.
 - ♦ Please let us know if that date would pose a significant issue.

SAM recalculation policy

- Delay for submission of GGUS tickets for recalculation should be respected. It is 10 days after monthly drafts of the availability reports are sent around by WLCG project office. If the GGUS tickets are not submitted respecting 10 days delay, they are not accepted;
- A/R recalculations will be accepted if they are relevant to the site MoU commitment or concern time ranges of sufficient length:
 - ♦ For T1 sites:
 - ◇ if only the corrected A/R will meet the 97% threshold;
 - ◇ or if the total concerned time ranges exceed 20 hours.
 - ♦ For T2 sites:
 - ◇ if only the corrected A/R will meet the 95% threshold;
 - ◇ or if the total concerned time ranges exceed 20 hours

Discussion

- Alberto: the 10-day timeline for recomputation requests needs to be respected.
- Maarten: normally it is, but there have been exceptions for various reasons, e.g. bugs.
- Maarten: CMS asked whether experiments still have flexibility in these matters. Experiments can indeed accept recomputation requests beyond the policy;

they can point to the policy when they do not want to accept a particular request.

- Alessandro: is this policy official now? In that case, could we follow it robotlike and could occasional glitches just be tolerated?
- Some discussion. It depends on the case (site, country, funding agency) to which extent glitches can be tolerated, but at least we now have a reasonable baseline policy.
- Stephan: CMS is used to doing recomputations shortly after relevant incidents, possibly long before the monthly reports are sent; in any particular case the affected number of hours could be a lot less than the given threshold (20); CMS is fine with that.
- Maarten: we will put the new policy on a dedicated Twiki page and explicitly clarify that the A/R machinery is always rerun for the final reports, irrespective of the number of recomputation requests (i.e. "silent" corrections are also picked up). Furthermore, we can change the policy in the future, should that be desirable.

Middleware News

- Useful Links
 - ◆ WLCGBaselineTable
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines/News
- Issues:
 - ◆ It was discovered that voms-clients-java-3.3.0 package in EPEL was broken. The update to that package caused essential commands (voms-proxy-info,init) to be removed. The investigation revealed that the cause of the problem was change of package name and the post-install scripts. Mattias Ellert provided a fixed package (voms-clients-java-3.3.0-2.el6) that is now in EPEL testing.

Discussion

- Balazs: the broken package was in EPEL-testing for weeks, but nobody reported anything; we should try to ensure EPEL-testing is checked at a few places.
- Maarten: agree. Let's follow up offline.

Tier 0 News

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- The activity levels have been lowish on average
 - ◆ Normal levels in the last days
 - ◆ New productions will be prepared
 - ◆ High amounts of analysis jobs in preparation of Quark Matter 2018 [🔗](#), May 13-19

- The 8th annual ALICE Tier-1/Tier-2 Workshop [📅](#) will take place next week April 17-19 in Derby, UK

ATLAS

- Stable grid production over the last weeks with up to ~300-350k concurrently running job slots, including the HLT farm. Additional HPC contributions with peaks of more than 1mio currently running job slots, but the actual corepower of the HPC CPUs is up to a factor of 10 weaker in comparison to regular grid site CPUs.
- Currently there is the usual mix of CPU heavy grid workflows of MC generation and simulation on-going with a smaller fraction of MC digitization and reconstruction. A smaller campaign of a delayed 2017 data stream was processed.
- Small operational hiccups due to certificate/proxy prolongations in various systems
- EOS reported potential corruption of files for an 8h period on 30 March - files may be corrupted even if Adler32 checksum is correct.
- Ongoing discussion within ATLAS to evaluate MD5 vs ADLER32 checksum.
- No operational problems with FTS in the past weeks, since large data reprocessing with large transfer rates finished. FTS at CERN and BNL in use at large scale and FTS at RAL in use for one site.
- Tier0 is ready for LHC data taking.

Discussion

- Alessandro: regarding the potentially corrupted files in EOS, in a similar incident 2 years ago there were 9k files affected - should we use MD5 instead of Adler-32? MD5 would be much safer, but cannot be calculated per stream when multiple streams are used in transferring files. If the file is still in memory after the Adler-32 calculation, the MD5 calculation only takes O(20s) more for a few GB.
- Christoph: CMS switched from CRC to Adler-32 a few years ago.
- Stephan: so far CMS did not see many corruptions while Adler-32 has been used.
- Alessandro: public clouds typically use MD5.
- Stephan: the best might be to have both in the experiment catalog and use the one that is best for a given storage. That may require code changes in quite some places.
- Christoph: that would be a major change.
- Maarten: would storage providers need to do something or is it just the experiments?
- Alessandro: for now ATLAS record both and use the MD5 sum only for debugging.
- Federico: in LHCb we did not have this discussion; it probably would be a big change, not a priority for now.
- Stephan: how many multi- vs. single-core jobs?
- Alessandro: ~100k single-core jobs; the rest of the "slots" should be divided by 8 to get an approximation of the number of multi-core jobs (the number of cores depends on the resource). We get a lot less performance out of HPC cores than their HS06 ratings suggest.
- Stephan: that could be due to the way KNL is used by the benchmark vs. real applications.

CMS

- cosmic data taken the last month with and without magnetic field
- LHC collisions since April but not physics data yet
- Had one round of transfer tests T0->T1
 - ♦ Might repeat a few tests after adjustments
- Compute systems busy at around 220k slots last month
 - ♦ usual 70% production 30% analysis split
- Kyungpook National University in Daegu, Korea informed us that they need to end Tier-2 service April 30th due to lost funding; KNU was an excellent site and we are sad to lose them
- Singularity deployment almost complete

- ◆ over 80% of Tier-1,2 sites ready
- ◆ one Tier-2 site and HLT still need to setup
- ◆ getting close to making Singularity mandatory in SAM (after ongoing installation activity at sites is complete)
- SAM corrections done as needed to make sure results/site evaluations are representative; do we need cross-experiment policy or can we leave it at VO discretion?
 - ◆ *See the SAM recalculation policy discussion above*

LHCb

- Ready to restart data taking
 - ◆ Use of HLTFarm for offline MC production has been close to 100% during the past few months, we expect reduction any time soon
- Productions:
 - ◆ Several Stripping productions close to the end. We could finish several productions since when CNAF is back in business. Run in "mesh" mode (~all Tier1s "helping" CNAF process its data)
 - ◆ MC simulation activities still taking close to 90% of distributed computing CPU
- CNAF:
 - ◆ almost all data recovered, including wet tapes. Staging for DataStripping went on without too many issues.
- DIRAC services
 - ◆ 6/7 weeks ago we experienced big problems with updating DBOD (MySQL and host): everything is operationally OK since then, but we'll still need to go through more upgrades ~soon
 - ◆ excluding the above, running 100% availability with an average of 120K concurrently running jobs
 - ◆ support for Glue2 is being added (late...): reports are that the same "mix of info" of BDII/Glue1 persists.
 - ◆ 1 new HPC site being integrated in these days.

Discussion

- Federico: GLUE-2 support had to be introduced because the CERN HTCondor CEs only publish that.
- Maarten: also the ARC CE will soon switch to GLUE-2 only. ALICE code will need to be adapted.
- Balazs: what rendering is used?
- Federico: LDAP.
- Alessandra: there is a proposal to switch to JSON instead.
- Julia: that would be OK for ATLAS and CMS, but ALICE and LHCb currently rely on dynamic CE info.
We will discuss it in the Information System Evolution TF meeting next week.
- Alessandra: we want that sites no longer have to patch their CE info providers at some point.

Ongoing Task Forces and Working Groups

Accounting TF

- Accounting Task Force meeting [in March](#) reviewed how experiments prepare the RRB reports.
- CERN batch accounting problem [is still under investigation](#)
- Julia: the standard Grafana plots are not optimal for certain reports.
- Alberto: Grafana supports downloading data in CSV format as input for other plot packages.
- Julia: a common solution would be nice. We will follow up.

Archival Storage WG

Update of providing tape info

Site	Info enabled	Plans	Comments
CERN	YES		
ASGC	NO		
BNL	YES		
CNAF	NO		
FNAL	YES		
IN2P3	NO		
JINR	NO		
KISTI	NO		
KIT	YES		
NDGF	NO		
NIKHEF-SARA	NO		
NRC-KI	YES		
PIC	YES		
RAL	NO		
TRIUMF	NO		

- Di: for what are those metrics needed?
- Julia: there is an effort to understand which metrics are useful to gain insight into the efficiency of tape operations; a basic set has been agreed and should be provided by all sites operating tape systems. Some metrics will be shown by the accounting portal, while others will go to a dedicated portal.
- Di: is it a manual process?
- Julia: no, the info should be provided automatically. Some sites already developed scripts that could be shared between sites with similar systems.
- Di: will follow up with our tape expert.
- Julia: the WG has a meeting on Monday April 16 at 16:00 CEST.
- Ron: will follow up for SARA.

Information System Evolution TF

- Next meeting is on the 19th of April [📅](#). Main topic of discussion is a proposal for alternative (vs BDII) way of publishing of the CE info.
- Alessandro: the CE info discussion does not only concern the BDII, but also the schema; we want it to be clear which resources are behind a particular CE.
- Federico: could that not be put into the GOCDB?
- Alessandro: yes. Mind that some development might be needed in the GOCDB.
- Alessandra: that is another point to be discussed next week.

IPv6 Validation and Deployment TF

Detailed status [here](#).

Machine/Job Features TF

Monitoring

NTR.

MW Readiness WG

Network Throughput WG

- perfSONAR 4.0.2 and CC7 campaign - 210 instances updated to 4.0.2; 81 instances already on CC7
 - ◆ **WLCG broadcast** will be sent to remind sites to plan an upgrade to CC7 and review the firewall port openings
 - ◆ perfSONAR 4.1 release, planned in Q2 2018 will no longer ship SL6 packages
- Attended perfSONAR developers F2F meeting in Amsterdam and presented feedback from OSG/WLCG
- WG reports planned for upcoming HEPiX and CHEP
- Networking and perfSONAR were also major topics at the OSG-All Hands (<https://indico.fnal.gov/event/15344/>)
 - ◆ 4 presentations were given on various topics related to the WG
 - ◆ One of the outcomes was a proposal to create a dedicated site-based documentation showing all links relevant to a given site
- WLCG/OSG network services
 - ◆ **Successfully migrated and commissioned new data pipeline**, we now have data flowing to UC/UNL/FNAL via RabbitMQ
 - ◆ Grafana was updated to reflect recent changes (<http://monit-grafana-open.cern.ch/dashboard/db/home?orgId=16>)
 - ◆ IPv6 dashboard [was added](#) to help compare IPv4 vs IPv6 performance side by side
- Outreach and other activities:
 - ◆ **GEANT has added several perfSONAR instances on LHCONE** at their major network hubs (ams, gva, lon, par, fra) - both IPv4 and IPv6
 - ◆ Advania was added to HNSciCloud test mesh
 - ◆ MGHPCC (<http://www.mghpcc.org/>) plans to deploy up to 22 perfSONARs, currently in discussion how we can help
- WLCG Network Throughput Support Unit: see twiki for summary of recent activities.

Squid Monitoring and HTTP Proxy Discovery TFs

Traceability WG

Container WG

Special topics

Action list

Creation date	Description	Responsible	Status	Comments
			DONE	

Creation date	Description	Responsible	Status	Comments
01 Sep 2016	Collect plans from sites to move to EL7	WLCG Operations		[older comments suppressed] Dec 7 update: Tier-1 plans are documented in the Nov 2 minutes. Jan 18 update: CREAM and the UI were released in UMD-4 on Dec 18. April 12 update: as various sites already upgraded OK, we no longer need to track this.
03 Nov 2016	Review VO ID Card documentation and make sure it is suitable for multicore	WLCG Operations	In progress	GGUS:133915 ↗
14 Sep 2017	Followup of CVMFS configuration changes, check effects on sites in Asia	WLCG Operations	CLOSED	March 1st update: this might imply significant effort; low priority for now. April 12 update: to be looked into when it is more urgent.

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

This topic: LCG > WLCGOpsMinutes180412

Topic revision: r14 - 2018-04-16 - 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