

Table of Contents

WLCG Operations Coordination Minutes, March 7, 2019.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Special topics.....	1
Follow up on EOS instabilities.....	1
Discussion.....	1
DPM upgrade, first experience and next steps.....	2
Discussion.....	2
Update of the WLCG Archival Storage Group.....	2
Discussion.....	2
Middleware News.....	2
Tier 0 News.....	3
Tier 1 Feedback.....	3
Tier 2 Feedback.....	3
Experiments Reports.....	3
ALICE.....	3
ATLAS.....	3
CMS.....	4
LHCb.....	4
Task Forces and Working Groups.....	4
GDPR and WLCG services.....	4
Accounting TF.....	4
Archival Storage WG.....	4
Update of providing tape info.....	4
Information System Evolution TF.....	5
IPv6 Validation and Deployment TF.....	5
Machine/Job Features TF.....	5
Monitoring.....	5
MW Readiness WG.....	5
Network Throughput WG.....	5
Squid Monitoring and HTTP Proxy Discovery TFs.....	6
Traceability WG.....	6
Container WG.....	6
Action list.....	6
Specific actions for experiments.....	6
Specific actions for sites.....	6
AOB.....	6

WLCG Operations Coordination Minutes, March 7, 2019

Highlights

- Follow up on EOS instabilities
- DPM upgrade, first experience and next steps
- Update of the WLCG Archival Storage Group

Agenda

<https://indico.cern.ch/event/803145/>

Attendance

- local: Mark Slater, Massimo Lamanna, Fabrizio Furano, Alessandra Doria, Giuseppe Bagliese, Oliver Keeble, Vlado Bahyl, Andrea Manzi, Concezio Bozzi, Gavin McCance, Julia Andreeva
- remote: Johannes Elmsheuser, Alessandro Paolini, Di Qing, Guillaume, Tim Bell, JM Barbet, Alessandro Cavalli, Dave Dykstra, Petr Vokac
- apologies:

Operations News

Special topics

Follow up on EOS instabilities

See the presentation [🔗](#)

Discussion

- Massimo presented the slides, explaining the reasons of the red areas of the availability plot. One red area is explained by ALICE instance catalogue upgrade. Both CMS and ATLAS issues reported at the previous meeting have been followed up with ATLAS and CMS. The catalogue upgrade planned for ATLAS and CMS has been agreed with the experiments.
- Upgrade for the version with bug fix (bug triggered by full replication groups) will be deployed after catalogue upgrade.
- Question about difference between FUSEx (which is currently running on EOSHOME) vs FUSE :
 - ◆ functionality is the same
 - ◆ better performance due to the local cache
 - ◆ administrative UI
 - ◆ still need to confirm stability, should be possible while running on EOSHOME
- Mark (LHCb) mentioned some instabilities of EOS recently experienced by LHCb. Not critical. Some operations failed, some outages short in time, not always SRM related. Massimo told that there were recently some EOS instabilities caused by non-expected network failures. The headnode was unreachable. Massimo asked Mark to send him the references to the opened tickets.

DPM upgrade, first experience and next steps

See the presentation [↗](#)

Discussion

- Julia asked whether it is possible to use different protocols including SRM by the same VO for different operations. Fabrizio told that it might inevitable in the close future, since VOs won't drop SRM quickly. The most critical problem of the parallel use of SRM along with other protocols for different operations is fixed in 1.12. The problem has been discovered by pioneer sites. Deletion with non-SRM protocol (http) has not been recognized by SRM which was used for writing
- Julia asked Alessandra and Petr representing pioneer sites their opinion about migration. Alessandra thinks that after 1.12 is out migration might be smooth, though it is difficult to foresee all problems which can be faced by sites which might migrate at the same time to the new OS
- Petr is already running 1.12 for two days. For the time being is fine, though it is very short time to make a conclusion. He also mentioned that there are several open tickets which are not critical for upgrade, however have to be addressed.
- Johannes expressed ATLAS concern regarding massive migration of ATLAS storage to the version which might bring troubles for experiment operations
- Petr also pointed out another important thing : Sites who do an upgrade should also enable gridftp redirection. In case of ATLAS, in order gridftp redirection to be used by ATLAS workflows, the change of site storage configuration in AGIS is required.
- What we do next:
 - ◆ Julia mentioned that ATLAS concerns are justified and should be taken into account. She suggested to wait for 1.12, then ask pioneer sites to deploy it, see for one month whether it does not cause problems and after a month select about 10 sites to take part in the second round of migration exercise.
 - ◆ People did agree with the plan. Oliver suggested to include in the second round diversity of sites , in terms of size, in terms of version they are currently running, whether they have or not puppet, etc...
- Julia asked Alessandro Paolini how WLCG should coordinate with EGI for the DPM upgrade campaign. Alessandro suggested to add him to the dpm-upgrade mailing list (done). At the point when we pass to massive migration the EGI broadcast will be sent

Update of the WLCG Archival Storage Group

See the presentation [↗](#)

Discussion

- Discussed the possibility to propagate/use dataset tag to/by the tape system. Dataset concept is used by all experiments. And dataset name is normally recorded in the namespace, so should be possible to recognize it by the tape system. However, it does not solve the problem of tape optimization since it is still unclear the combination of datasets which will be recalled simultaneously.
- Vlado and Oliver stressed the importance of the experiments to take part in the work of the group. Experiments in the group are currently mostly represented by ATLAS and CMS. LHCb has been invited to take part. LHCb tended to agree.

Middleware News

- Useful Links
 - ◆ WLCGBaselineTable
 - ◆ MW Issues
 - ◆ Storage Deployment

- Baselines/News

Tier 0 News

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- High activity on average
- No major issues
- CERN
 - ◆ EOS: smooth migration to QuarkDB back-end
 - ◆ HTCondor issue:
 - ◇ random CEs develop very large numbers of waiting jobs, while other CEs could run them
 - ◇ the VOBOXes stop submitting jobs when the total number of waiting jobs exceeds a threshold
 - ◇ we have regularly needed to exclude the "stuck" CEs, to avoid losing part of the ALICE fair share
 - ◇ and put them back when their backlog has been handled
 - ◇ such tuning activities were not needed last year
 - ◇ we will look into making the CE selection smarter
 - now it is round-robin

ATLAS

- Smooth Grid production over the last weeks with ~300k concurrently running grid job slots with the usual mix of MC generation, simulation, reconstruction, derivation production and analysis and a small fraction of dedicated data reprocessing. Some periods of additional HPC contributions with peaks of ~150k concurrently running job slots last month and ~15k jobs from Boinc. The HLT farm/Sim@P1 was available for one week adding ~100k additional job slots.
- Commissioning of the Harvester submission system via PanDA is in its last steps: most grid clouds have been migrated to unified queues using Harvester.
- Started the gradual commissioning of a new PanDA worker node pilot version in production. This process will last for the next months and will replace completely the existing pilot.
- ATLAS sites jamboree and HPC strategy meeting, 5-8 March at CERN, <https://indico.cern.ch/event/770307/>
- Reminders still valid (already advertised, small updates):
 - ◆ CentOS7: ATLAS would like to start a more forceful migration to CentOS7 and have the vast majority of resources, if not all, migrated by June 1.
 - ◆ SCRATCHDISK: ATLAS would like to increase the SCRATCHDISK quota to 100TB per 1000 analysis slots
 - ◆ IPv6: if sites update to IPv6 dual-stack please let us know in advance, SAM tests have been developed

CMS

- smooth running, compute systems busy at about 230k cores
 - ◆ usual production/analysis mix (75%/25%)
- network overload at In2P3 did not impact CMS significantly the last weeks
- 2017 and 2018 Monte Carlo production ongoing
- preparing (small) tape deletion campaign
- no severe EOS issue last month (i.e. still some problem, but no service interruption)
- CRAB/CMS user batch service hit CPU stealing on hypervisors; OpenStack team looking into ways to mitigate this;
- CMSR database running queries slowly due to lost Oracle optimization when we switched user stage out (two different queries simultaneously)

LHCb

- MC simulation and user analysis with ~50K jobs running
- Main issues to report are the EOS instabilities. Background of read/write failures with occasional peaks of very little access at all
- Need to agree on who's going to do the re-computation for SAM test failure in last week.
 - ◆ The question was mostly what to do with the time period where no test were run. Agreed to put it to 'unknown', which won't have impact on re-calculation. Re-calculation has been performed by the monitoring team the day after the meeting.

Task Forces and Working Groups

GDPR and WLCG services

- Updated list of services
- Julia mentioned that the first draft of the template of the privacy note is ready and is currently under discussion.

Privacy note should be enabled by all services dealing with user data

Accounting TF

- First storage space accounting reports (non-official for the time being) for January have been sent around by the WLCG project office. Sites have been kindly asked to look into the numbers and report problems/inconsistencies. The WSSA developers already got feedback from some sites. Julia expressed gratitude to the sites for their input, it helps a lot to find issues with the system
- Accounting status/plans and data validation will be presented at the GDB next week

Archival Storage WG

Update of providing tape info

PLEASE CHECK AND UPDATE THIS TABLE

Site	Info enabled	Plans	Comments
CERN	YES		
BNL	YES		
CNAF	YES		Space accounting info is integrated in the portal. Other metrics are on the way

Site	Info enabled	Plans	Comments
FNAL	YES		
IN2P3	YES		Space accounting info is integrated in the portal. Other metrics are on the way
JINR	YES		
KISTI	YES		KISTI has been contacted. Will work on in the second half of September
KIT	YES		
NDGF	NO		NDGF has a distributed storage which complicates the task. Discuss with NDGF possibility to do aggregation on the storage space accounting server side. Should be accomplished by the end of the year
NLT1	YES		Almost done, waiting for opening of the firewall, order of couple of days
NRC-KI	YES		
PIC	YES		Space accounting info is integrated in the portal. Other metrics are on the way
RAL	YES		Space accounting info is integrated in the portal. Other metrics are on the way
TRIUMF	YES		

One can see all sites integrated in storage space accounting for tapes [here](#)

Information System Evolution TF

- NTR

IPv6 Validation and Deployment TF

Detailed status [here](#).

Machine/Job Features TF

Monitoring

MW Readiness WG

Network Throughput WG

- **perfSONAR infrastructure status - CC7/4.1 campaign ongoing**
 - ◆ perfSONAR 4.0 and perfSONARs on SL6 are no longer supported since Q4 2018 - please update ASAP
 - ◆ **New baseline version for perfSONAR is the latest release 4.1.6 (fixes important bug causing duplicate testing)**
- WLCG/OSG network services were updated
 - ◆ All meshes were updated to test throughput and traceroutes over both IPv4 and IPv6; dual stack mesh was retired
 - ◆ Monitoring was updated with new thresholds and now also tracks IPv4/IPv6 efficiency (https://psetf.aglt2.org/etf/check_mk/)
 - ◆ Documentation was updated as well (<https://opensciencegrid.org/networking/>)
 - ◆ **perfSONAR dashboard ([psmad/maddash](#)) was re-configured and fixed**
- WLCG Network Throughput Support Unit: see twiki for summary of recent activities.

Squid Monitoring and HTTP Proxy Discovery TFs

- Nothing to report
- Discussion about recommended squid version will take place at the next meeting

Traceability WG

Container WG

Action list

Creation date	Description	Responsible	Status	Comments
---------------	-------------	-------------	--------	----------

Specific actions for experiments

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

Specific actions for sites

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

AOB

- If there is no major objection, next meeting would take place the 4th of April. We will have WLCG workshop in JLAB in between. So we might discuss from the outcome of the workshop what concerns WLCG operations at the April meeting.

This topic: LCG > WLCGOpsMinutes190307

Topic revision: r13 - 2019-08-21 - JuliaAndreeva



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