

Table of Contents

WLCG Operations Coordination Minutes, Dec 6, 2018.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Special topics.....	1
IPv6 deployment update.....	1
Discussion.....	1
Middleware News.....	2
Discussion.....	2
Tier 0 News.....	2
CC7.....	2
Discussion.....	3
LSF decommissioning.....	3
Tier 1 Feedback.....	3
Tier 2 Feedback.....	3
Experiments Reports.....	3
ALICE.....	3
ATLAS.....	4
CMS.....	4
LHCb.....	4
Task Forces and Working Groups.....	4
GDPR and WLCG services.....	5
Accounting TF.....	5
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.....	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
Discussion.....	6
Action list.....	6
Specific actions for experiments.....	6
Specific actions for sites.....	6
AOB.....	6

WLCG Operations Coordination Minutes, Dec 6, 2018

Highlights

- IPv6 deployment update
- CentOS / CC7 timelines for lxbatch and lxplus
- Container WG update
- **CREAM CE end of support is Dec 2020**

Agenda

<https://indico.cern.ch/event/778285/>

Attendance

- local: Andrea (WLCG), Christoph (CMS), Gavin (T0), Giuseppe (CMS), Maarten (ALICE + WLCG), Rob (LHCb)
- remote: Andreas (KIT), Catherine (LPSC + IN2P3), Eric (IN2P3-CC), Felix (ASGC), Gareth (RAL), Johannes (ATLAS), Marcelo (CNAF), Matt (EGI), Miro (WLCG), Stephan (CMS), Thomas (DESY)
- apologies:

Operations News

- the next meeting is planned for *Feb 7*
 - ◆ please let us know if that date would pose a significant problem

Special topics

IPv6 deployment update

see the presentation [↗](#)

Discussion

- should the deadline be extended?
- what would be the consequences for sites that do not meet the deadline?
- Stephan: could we think of a "carrot" for sites?
 - ◆ for example, a future CERN cloud extension might only offer IPv6
 - ◆ only experiments that are ready would be able to use those resources
 - ◆ experiments would have an extra reason to put pressure on their sites
- Maarten: some sites may be unable to move faster, whatever the pressure
- Andrea: some sites may even depend on external factors
 - ◆ Tokyo want to ascertain that routes e.g. to Canada work OK via IPv6
 - ◇ no response yet from CANARIE

- Maarten: in principle experiments could adjust their frameworks and workflows to deal with a fraction of their sites not being ready, but such efforts may be expensive and hence are best avoided altogether
 - ◆ possibly doable in the end for just a small tail of sites
- Christoph: it would otherwise be a major development that we have to avoid

- Stephan: IPv6 readiness can also be lost again if it is not exercised
 - ◆ we already discovered some cases
- Andrea: we plan to have IPv6-only additional SAM instances in the near future
- Stephan: that will give us a "stick"
 - ◆ we can mark the failing sites with a warning state and color
 - ◆ the A/R calculations would not be affected for the time being
- Maarten: in the future such tests could even become critical
- Catherine: that would be too strong for the time being

- Maarten:
 - ◆ so far the progress looks quite good after all
 - ◆ next year we should have IPv6 as a standing item on the agenda
 - ◆ the IPv6-only additional SAM instances should be set up early next year
 - ◆ at some point we may conclude if another deadline makes sense or not

Middleware News

- Useful Links
 - ◆ WLCGBaselineTable
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines/News
 - ◆ **CREAM CE end of support is Dec 2020**
 - ◆ Recently discovered top BDII issues are fixed in UMD 4.8.0
 - ◇ Stale list of site BDIIs due to recent python versions no longer permitting transparent redirection from *http* to *https* GOCDB URL
 - ◇ Continued attempts to obtain information about OSG services long after the OSG BDII has been retired

Discussion

- Matt: we should look together into a strategy for phasing out CREAM
- Maarten: yes
- Marcelo: could WLCG advise on a replacement?
- Maarten:
 - ◆ the ARC CE and the HTCondor CE are both viable replacements
 - ◆ the HTCondor CE should soon have proper APEL accounting support
 - ◆ sites will need to check with their other customers
 - ◆ there is plenty of time, but sites should already consider these matters

Tier 0 News

CC7

The schedule for the upgrade of CERN-PROD to CC7 has now been published (OTG:0047300 [↗](#)).

We will begin the migration of the public CERN Batch Service resources to CC7, starting in January 2019, according to the following schedule:

- 30% of the public and grid batch resources will be available as CC7 by the end of January 2019
- 50% of the public and grid batch resources will be available as CC7 by the end of March 2019
- On the 2nd of April:
 - ◆ The default submit target for HTCondor will be changed to CC7.
 - ◆ The lxplus.cern.ch alias will switch to point to the CC7 LxPlus service.
- The remainder of Batch capacity will be migrated over to CC7 by early June 2019.

Note that the old SLC6 LxPlus service will remain accessible at lxplus6.cern.ch.

(The CC7 based LxPlus service can be accessed in advance of the change at lxplus7.cern.ch)

For selecting the non-default batch operating system for HTCondor job submission, please see CERN Batch Service user guide [?](#).

Note that CERN dedicated resource pools will be handled separately with the relevant groups and experiment teams (e.g, Tier-0, CMSCAF, ...)

Singularity is available now for (power) users to swap SLC6 / CC7 (as per CMS production) - we're looking at ways of making this more transparent.

Discussion

- Maarten: do these plans look OK to the experiments?
- Christoph:
 - ◆ OK for CMS
 - ◆ we particularly look forward to an easy way to let jobs ask for SLC6
- Rob: OK for LHCb
- Johannes: OK for ATLAS
- Maarten: looks OK also for ALICE

LSF decommissioning

Reminder: OTG:0046088 [?](#) - the LSF public share will be switched off at end January 2019. All remaining local users of the public share should move to HTCondor. Grid users have already moved.

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- Successful end of the Run 2 data taking!
 - ◆ HI raw data replication ongoing
 - ◆ T1 sites mostly done

◇ Thanks to expert teams at CERN and T1 sites!

- ◆ Reconstruction and reprocessing for many months to come
- ◆ Plus lots of analysis and MC simulation as usual

ATLAS

- Successful end of Run2 data taking. The data transfers of the Heavy Ion data from CERN Point1 to EOS to Tape and the Tier1s performed well. The only (predicted) bottleneck occurred in transfers from EOS to Castor disk at 2.5-3 GB/s. This bottleneck was mitigated by temporarily not writing Tier0 AODs to tape at CERN.
- Smooth Grid production over the last weeks with ~300k concurrently running grid job slots. The heavy ion data is processed at the Tier0 with its full capacity and about 20-30k Grid job slots are used in spillover mode. After the finish of Run2 about 90k jobs slots from the HLT farm are now used with Sim@P1. Additional HPC contributions with peaks of ~100k concurrently running job slots last month and ~11k jobs from Boinc.
- Commissioning of the Harvester submission system via PanDA is on-going: most grid clouds have been migrated to unified queues using Harvester. The only remaining cloud is the US. In January non-unifiable queues and analysis queues are planned to move to Harvester submission.
- IPv6: Since one week network problems between aCT (ARC control tower) and PanDA servers (INC1859087 [↗](#))

CMS

- data taking finished successfully
 - ◆ heavy ion data volume a bit less than anticipated
 - ◆ half of the heavy data already transferred to the secondary archive site, Fermilab
 - ◆ transfer and reconstruction at CERN expected to continue through December
- several issues with SAM infrastructure
 - ◆ SAM ETF hanging on CEs three weeks ago, thanks Marian for clearing/restarting!
 - ◆ SAM3 not updating two weekends ago and again last weekend, Thursday to Monday outage in each case
 - ◇ last one self-inflicted due to last globus toolkit CE being removed from service in October, thanks for identifying this Luca!
- no EOS crash but
 - ◆ serious EOS checksum errors on last day of data taking INC:1855937 [↗](#)
 - ◆ data server overload GGUS:138506 [↗](#)
- Monte Carlo production for 2018 configuration and re-reconstruction of early 2018 data in progress
- compute systems busy at about 200k cores, usual mix of about 75% production and 25% analysis
- storage IPv6 accessible at about 60% of CMS sites and almost or partially IPv6 accessible at another 7%

LHCb

- Successful end of the Run 2
- At the start of next year there will be a new stripping campaign on Run1 and Run2 data, staging has already begun for this.
- Next year there will also be an emphasis on MC production.
- Continuing to have IPv6 problems with data transfers at SARA with multiple GGUS tickets open

Task Forces and Working Groups

GDPR and WLCG services

- Updated list of services

Accounting TF

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
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 [here](#)

Information System Evolution TF

IPv6 Validation and Deployment TF

Detailed status here.

Machine/Job Features TF

Monitoring

MW Readiness WG

Network Throughput WG

Squid Monitoring and HTTP Proxy Discovery TFs

- Nothing to report

Traceability WG

Container WG

See WLCGContainers for details and last meeting [↗](#).

- ATLAS and CMS presented their usage of containers - the working group noted the interest for ATLAS for (potentially many) user-level analysis containers, which will likely form the basis of ATLAS' analysis reproducibility
- CERN EP presented the "unpacked.cern.ch" service in development which takes a docker URL and unpacks the directory tree in CVMFS (as per "singularity.opensciencegrid.org") and additionally can provide CVMFS-hosted layers to docker clients using the CVMFS-driver. Experiments requested interest to test as soon as its available.
- Future topics to discuss:
 - ◆ The release of CC 7.6 (with unprivileged user namespaces) means experiments and sites could run Singularity in unprivileged mode in production
 - ◆ Fat images vs. CVMFS-unpacked images

Discussion

- Maarten:
 - ◆ we probably need to discuss the plans for user containers further
 - ◆ not clear how the CVMFS cache turnover rates might be affected

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

This topic: LCG > WLCGOpsMinutes181206

Topic revision: r15 - 2018-12-10 - MaartenLitmaath



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback