

Table of Contents

WLCG Operations Coordination Minutes, November 8th, 2018.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Special topics.....	1
Hammer Cloud for commissioning of compute resources.....	1
Discussion.....	1
Middleware News.....	2
Tier 0 News.....	2
Tier 1 Feedback.....	2
Tier 2 Feedback.....	2
Experiments Reports.....	2
ALICE.....	2
ATLAS.....	2
CMS.....	3
Discussion on lxbatch CentOS 7 migration.....	3
LHCb.....	4
Task Forces and Working Groups.....	4
GDPR and WLCG services.....	4
Accounting TF.....	4
Discussion.....	4
Archival Storage WG.....	4
Update of providing tape info.....	4
Information System Evolution TF.....	5
IPv6 Validation and Deployment TF.....	5
Discussion.....	5
Machine/Job Features TF.....	5
Monitoring.....	5
MW Readiness WG.....	5
Network Throughput WG.....	5
Squid Monitoring and HTTP Proxy Discovery TFs.....	6
Discussion on Squid Monitoring.....	6
Discussion on HTTP Proxy Discovery.....	7
Traceability WG.....	7
Container WG.....	7
Action list.....	7
Specific actions for experiments.....	7
Specific actions for sites.....	7
AOB.....	7

WLCG Operations Coordination Minutes, November 8th, 2018

Highlights

- Hammer Cloud for commissioning of compute resources
- Discussion on Ixbatch CentOS 7 migration
- Sites should upgrade their perfSONAR services to v4.1 on CentOS 7
- Discussion on Squid Monitoring
 - ◆ Operations experts being sought

Agenda

<https://indico.cern.ch/event/769760/>

Attendance

- local: Alessandro (ATLAS), Ben (CERN computing), Jarka (CERN computing), Julia (WLCG), Maarten (ALICE + WLCG), Marian (networks + monitoring)
- remote: Christoph (DESY), Darren (RAL), Dave D (FNAL), Di (TRIUMF), Felix (ASGC), Gareth (RAL), Javier (IFIC), Jeremy (GridPP), Stephan (CMS), Thomas (DESY)
- apologies:

Operations News

- The next meeting is planned for Dec 6
 - ◆ Please let us know if that date would pose a significant problem

Special topics

Hammer Cloud for commissioning of compute resources

see the presentation

Discussion

- Julia: can you test batch resources without an intermediate job submission layer?
- Jarka: indeed
- Maarten: you would at least need to get an experiment proxy?
- Alessandro, Jarka:
 - ◆ right, but from then on HC can test resources independently
 - ◆ also the HC monitoring is self-contained
- Julia:
 - ◆ HC allows many involved services to be tested
 - ◆ it could be an interesting opportunity for other sites
- Maarten:
 - ◆ might the system make a CERN-specific assumption somewhere?
 - ◆ can other sites download something and try it out?

- Jarka:
 - ◆ for now it only works for HTCondor CE instances, but we can help with porting it to other CE types
 - ◆ it is not yet packaged; we first would like to gauge the interest of other sites
- Julia: for other sites to get it to work, there is a dependency on experiments?
- Jarka: indeed, and the configuration can be tricky, but we can help with that
- Christoph: will the current CRAB workflows stay?
- Jarka: yes, the usage presented today is separate

Middleware News

- Useful Links
 - ◆ WLCGBaselineTable
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines/News
 - ◆ An EGI broadcast [↗](#) concerning TLS v1.2 support on EGI / WLCG has been sent to all EGI sites on Oct 29

Tier 0 News

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- Normal activity levels on average
- No major issues on the grid
- Task queue was affected by a HW I/O problem late Oct - early Nov
 - ◆ Causing erratic activity levels and job failures
 - ◆ Resolved by replacing the machine

ATLAS

- Smooth Grid production over the last weeks with ~330k concurrently running grid job slots.
 - ◆ During last week's LHC TS and MD additional ~90k jobs slots from the HLT farm with Sim@P1. Such a big single resource is interesting for ATLAS all together, since it's like a stress test for parts of the central ADC infrastructure.
 - ◆ Additional HPC contributions with peaks of ~100k concurrently running job slots and ~10k jobs from Boinc.
- Commissioning of the Harvester submission system via PanDA is on-going: currently finishing the migration of the UK cloud and starting with the remaining clouds in FR, NL, CA, DE and the US.
- The Heavy Ion data throughput from CERN Point1 to EOS to Tape will be higher than initially planned and used in the September throughput test. I.e. 4.5GB/s during the run. Still 50% duty cycle expected (from LHC, not up to ATLAS to decide). This should be just "correct" with the tested

throughput to tape (2-2.5GB/s): backup plans in case of troubles have been defined.

- Asked sites to slowly migrate to CentOS7 in the coming months
- Crash of EOSATLAS on Wednesday morning: [link](#)

CMS

- Setup for HI run ongoing
- Good CPU utilization recently
 - ◆ ~160k cores for production and ~50k cores for analysis
- Emergency stop of production activities ~two weeks ago because of disk shortage
 - ◆ We reached 85% of "unmovable data"
 - ◆ Situation improved (back to 68% of "unmovable data") thanks to some cleaning performed and to early availability of 2019 pledges from some sites
- Still some EOS instabilities related to fuse mount INC:1784940 (recent EOS crashes OTG:0046403, OTG:0046125)
 - ◆ Getting a quite good support, with daily interactions and fixes
- Upgrading several services to CentOS7
- Planned switchover to CentOS7 for CMSSW 2019 releases
 - ◆ Singularity makes the change transparent for most of the MC and Analysis jobs submitted through grid
 - ◆ Still need solutions with sl6/sl7 availability for users running "local job". We proposed IT to look into a solution which uses singularity at the level of batch system, we would like their feedback on that

Discussion on lxbatch CentOS 7 migration

- Christoph:
 - ◆ it would be good if our users could use a Condor classad to have their jobs run in a container with the required OS
- Stephan:
 - ◆ we would like lxbatch to handle SLC6 and CC7 more dynamically
 - ◆ at some point there could be a peak demand for SLC6, e.g. shortly before a conference: how can that be handled?
- Ben:
 - ◆ the initial plan was to start with 10% CC7 in Jan, 50% (80%) in June 2019 and 100% in June 2020 or so
 - ◆ we would prefer moving faster, but only if it is OK for the experiments
 - ◆ users can choose the required OS now
 - ◆ the default currently is SLC6 because the majority of lxbatch has that OS
 - ◆ when the latter shifts to CC7, we adjust the default accordingly
 - ◆ and let the lxplus alias point to lxplus7
- Stephan: could the job OS be taken from the lxplus version?
- Ben:
 - ◆ that cannot be done easily; we prefer the user specify it explicitly
 - ◆ see <http://batchdocs.web.cern.ch/batchdocs/>
- Ben:
 - ◆ CC7 comes with increased capabilities to use containers:
 - ◆ more support in CVMFS
 - ◆ HTCondor Docker universe

- Stephan:
 - ◆ users should not need to do Singularity or Docker manipulations
 - ◆ can a classad be made sufficient to get an SLC6 environment?
- Ben: when only CC7 is left, we can use containers internally to still satisfy SLC6 requirements
- Alessandro:
 - ◆ this matter depends very much on *when* the number of SLC6 (CC7) resources is going down (up)
 - ◆ by the time SLC6 is at 1% there could be an issue for SLC6 users
- Ben: by then we could apply some magic as needed
- Alessandro:
 - ◆ a long co-existence of the 2 OS is a bit painful for ATLAS
 - ◆ could that magic already be applied now?
- Ben:
 - ◆ the big issue with containers is the support of AFS
 - ◆ also for us it is undesirable to maintain 2 *large* pools in parallel
- Stephan: the ramp-down of SLC6 is hard to predict
- Maarten: in the coming months we will keep revisiting this matter, to see if we can do the transition faster

LHCb

Task Forces and Working Groups

GDPR and WLCG services

- Updated list of services

Accounting TF

- The latest WLCG Accounting Task Force discussed HTCondor accounting. We were lacking standard solution for configuration HTCondorCE + HTCondor batch. Stephen Jones evaluated PIC implementation and is evolving it further so that it would be straight forward to use it by all sites which have such configuration.

Discussion

- Ben: got contacted by IHEP Beijing about HTCondor accounting
- Julia: they would exactly need the solution we are working towards
- Ben: the HTCondor devs offered help with accounting issues
- Julia: for now the matter lies rather with APEL

Archival Storage WG

Update of providing tape info

PLEASE CHECK AND UPDATE THIS TABLE

Site	Plans	Comments
------	-------	----------

	Info enabled		
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

- Work is specification of the computing resource description (Computing resource Reporting- CRR) in order to provide an alternative of the CE description via BDII.

IPv6 Validation and Deployment TF

Detailed status here.

Discussion

- Jeremy (in the chat window):
 - ◆ it looks unlikely that most T2 sites will have their storage dual-stack by the end of Run 2
- Maarten:
 - ◆ the end of Run 2 has always been an approximate timing
 - ◆ early 2019 would already be great
- Julia: we will ask for an update in the next meeting

Machine/Job Features TF

Monitoring

MW Readiness WG

Network Throughput WG

- **perfSONAR infrastructure status - CC7/4.1 campaign ongoing**
 - ◆ Sites were reminded to upgrade to CC7 and review their configuration (preferably by end of October)

- ◆ Still only around 50% of nodes are on CC7 as of today - we'll soon start contacting sites directly
- ◆ Some sites waiting for/deploying new hardware; e.g. SARA deployed 100Gbps perfSONAR (first in Europe), BNL deployed 2x40 Gbps perfSONAR
- **WG update was presented at HEPiX and LHCOPN/LHCONE workshop**
 - ◆ LHCOPN/LHCONE workshop had a dedicated talk on MTU - topic that was raised by the WG - draft recommendation was proposed (<https://indico.cern.ch/event/725706/>)
- WLCG/OSG network services working fine
- **WLCG Network Throughput Support Unit:** see twiki for summary of recent activities.
 - ◆ Number of cases recently popped up, some are still unresolved and are being followed up by the corresponding R&Es and/or will be raised at the upcoming Asia Tier Forum
 - ◆ Case in point (NCP):
<https://docs.google.com/document/d/1HHhK9t4PpYPzZOofJUAhupRAodhPT6HNIZi6J8ljpBtw/edit#>

Squid Monitoring and HTTP Proxy Discovery TFs

- Especially now that we have the CVMFS failover monitor [?](#), there is a need for a WLCG operations person who watches it and helps sites whose clients are failing to use their own squids
 - ◆ Development staff has been beginning to do this, but that is not sustainable
 - ◆ CMS has had a person to do this (Barry Blumenfeld) for many years for Frontier
 - ◆ ATLAS needs someone as well
 - ◆ It would make sense for one person to do it for all squid failovers
- A cernvm-wpad service is being readied for production use at CERN and FNAL for the purpose of making the default CernVM proxy configuration work out of the box anywhere
 - ◆ All traffic will go through openhtc.io Cloudflare aliases
 - ◆ The concern then is that it may be used in large numbers in cloud deployments without their own squids
 - ◇ May use Cloudflare too heavily, and also worse performance than local network squids
 - ◆ The new cernvm-wpad service will use the following algorithm:
 1. Return WLCG squids if at a WLCG site. In the future that will include cloud squids registered via Shoal
 2. Otherwise track the request rates per Geo IP organization, and if they're coming in large volumes from any organization, direct traffic through backup proxies which will be monitored
 3. Otherwise return DIRECT to go through Cloudflare

Discussion on Squid Monitoring

- Julia: if an issue is detected automatically, the TF still helps?
- Dave: it should be an ops person who gets involved, not a dev

- Julia: do the experiments notice such issues in their monitoring?
- Alessandro: yes, and they open tickets for the affected sites
- Dave: a shifter can notice a problem, but not solve it

- Alessandro:
 - ◆ the problem can be due to the workflow, not the squid
 - ◆ e.g. overlay jobs

- Julia: should we rather have a support unit in GGUS instead of relying on someone monitoring for errors?
- Dave: we would like to have an intermediary person that would be able to help sites solve most of the issues

- Maarten:
 - ◆ as ATLAS and CMS are the biggest users, it would seem reasonable that such intermediary persons would be drawn from these experiments
 - ◆ it is unrealistic to expect an ALICE or LHCb person to work on debugging workflow issues in ATLAS or CMS
 - ◆ as the WLCG ops team is very small these days, we must be very careful with committing effort there
- Dave: it would be great to have 1 person from each of ATLAS and CMS
- Julia: we will follow up offline

Discussion on HTTP Proxy Discovery

- Maarten:
 - ◆ the algorithm looks OK
 - ◆ it would be good to keep track of how much Cloudflare is used, to try and avoid a surprise blocking of our clients
 - ◆ let's anyway give it a go and see what happens

Traceability WG

Container WG

- Dave: there are 3 highlights concerning Singularity:
 - ◆ mount points now can be added to containers that do not already have the target directories in their image
 - ◇ that is important for ATLAS and also CMS want to use it
 - ◇ the feature is called `underlay` and is available as of version 2.6
 - ◇ in that version it is not enabled by default
 - ◆ RHEL 7.6 (released Oct 30) supports unprivileged mount namespaces
 - ◇ that allows Singularity to be run *unprivileged* with *underlay* instead of *setuid* with *overlay*
 - ◆ version 3.0.1 has been released by the devs
 - ◇ written in Go
 - ◇ should be compatible with 2.6
 - ◇ new features that we do not need at this time
 - ◇ being tested

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

-- JuliaAndreeva - 2018-10-08

This topic: LCG > WLCGOpsMinutes181108

Topic revision: r26 - 2018-11-12 - MaartenLitmaath



Copyright &© 2008-2021 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](#)