

Table of Contents

WLCG Operations Coordination Minutes, March 2th 2017.....	1
Highlights.....	1
Agenda.....	1
Attendance.....	1
Operations News.....	1
Middleware News.....	2
Tier 0 News.....	2
Tier 1 Feedback.....	3
Tier 2 Feedback.....	3
Experiments Reports.....	3
ALICE.....	3
ATLAS.....	3
CMS.....	3
LHCb.....	4
Discussion.....	4
Ongoing Task Forces and Working Groups.....	4
Accounting TF.....	4
Information System Evolution TF.....	4
IPv6 Validation and Deployment TF.....	4
Machine/Job Features TF.....	5
Monitoring.....	5
MW Readiness WG.....	5
Network and Transfer Metrics WG.....	5
Squid Monitoring and HTTP Proxy Discovery TFs.....	5
Traceability and Isolation WG.....	5
Theme: SLAs and usage of different kinds of computing resources.....	6
Discussion.....	6
Theme: Machine/Job Features update.....	6
Discussion.....	6
Action list.....	7
Specific actions for experiments.....	8
Specific actions for sites.....	8
AOB.....	8

WLCG Operations Coordination Minutes, March 2th 2017

Highlights

- Lessons learned from ATLAS tape usage tests. Follow up via FTS steering group. Share experience with other experiments.
- "Mid SLA" resources. During the next months accumulate experience at CERN (providing such resources) and in experiments (using such resources). Review accumulated experience after a couple of months and decide what we need to do in terms of tagging, accounting, etc...
- MJF deployment is not as successful as desired. Find help (at least temporary) for deployment campaign concentrating on the LHCb sites. See whether this task would become more important considering outcome of the benchmarking discussion.

Agenda

- <https://indico.cern.ch/event/615987/>

Attendance

- local: Alessandro (ATLAS), Andrea M (MW Officer + data management), Andrew (LHCb + Manchester), Gavin (T0), Julia (WLCG), Maarten (WLCG + ALICE), Marcelo (LHCb), Vincent (security)
- remote: Alessandra D (Napoli), Antonio (PIC), Catherine (IN2P3 + LPSC), Christoph (CMS), CNAF, David M (FNAL), Di (TRIUMF), Felix (ASGC), Javier (IFIC), Kyle (OSG), Leonardo (Sussex), Renaud (IN2P3-CC), Ron (NLT1), Stephan (CMS), Vincenzo (EGI), Xin (BNL)

Operations News

- This year's WLCG workshop will be held June 19-22 in Manchester
 - ◆ Agenda: <http://indico.cern.ch/event/609911/>
 - ◆ Feel free to suggest topics, thanks! ([wlcg-ops-coord-chairpeople \[at\] cern.ch](mailto:wlcg-ops-coord-chairpeople[at]cern.ch))
- The WLCG Data Management steering group had 2 kick-off meetings
 - ◆ The mandate, list of tasks and priorities are being finalized
 - ◇ *Julia: for example, storage space accounting is in the list of topics*
- PIC is developing an APEL parser for HTCondor
 - ◆ Easily adaptable by other sites
 - ◆ More in the next meeting
- In order to avoid inconsistencies in the naming of the service types along the WLCG IS chain (GocDB, OIM, SAM, experiment-specific systems like Dirac), agreed with representatives from OSG, EGI and GocDB team and IS evolution task force on the policy for introducing of the new service types. The request for the introducing of the new service types should be sent to is-approvals@cernNOSPAMPLEASE.ch list. The name for a new service type will be agreed among the members of this list and then can be introduced in GocDB , OIM, etc... Information about new service types will be then broadcasted to the experiments and members of the IS evolution task force which includes members who can be concerned as for example members of the monitoring team.
- The next Ops Coordination meeting will be on April 6

Middleware News

- Useful Links:
 - ◆ [Baseline Versions](#)
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines/News:
 - ◆ Baselines updated: removed dCache 2.10, moved dCache 2.13 baselines to 2.13.51 which fixes an issue with RFC proxy for certain CAs and improve bulk deletions, FTS moved to 3.5.7
 - ◆ dCache 2.10.x support ended in 2016. We discussed with EGI and prepared a broadcast together (already sent), still 16 instances running this version. EGI will open tickets soon.
- Issues:
 - ◆ High risk CVE-2017-6074 Linux kernel privilege escalation vulnerability (<https://wiki.egi.eu/wiki/SVG:Advisory-SVG-CVE-2017-6074>). Sites should apply the kernel patches or the mitigations as reported in the advisory.
 - ◆ 2 issues discovered in the latest Xrootd release (4.6.0) both client and server side. Sites/Experiments are suggested not to upgrade to this version and wait for 4.6.1 under preparation.
- T0 and T1 services
 - ◆ ASGC
 - ◇ DPM upgrade to v 1.8.11
 - ◆ BNL
 - ◇ Enabled dual stack FTS
 - ◆ CERN
 - ◇ check T0 report
 - ◇ FTS upgrade to v. 3.5.8 and gfal2 2.13.1
 - ◆ IN2P3
 - ◇ Migration of the core dCache servers to CentOS7 (postgres 9.5, dcache 2.13.54)
 - ◆ JINR
 - ◇ dCache minor upgrade 2.13.51 -> 2.13.54, Postgres minor upgrade 9.4.9 -> 9.5.1
 - ◆ KIT
 - ◇ FAX decommissioned and dCache updated to 2.13.51 for ATLAS on 1st Feb
 - ◆ NL-T1
 - ◇ SURFsara upgraded dCache from 2.13.49 to 2.13.51 on Feb 2.
 - ◆ PIC
 - ◇ Enstore upgraded, dCache upgrade planned for March 3rd
 - ◆ RAL
 - ◇ Castor 2.1.15-20 update recently completed. All data now on T10KD dives/media.
 - ◇ gfal2 upgraded to v 2.13.1 on FTS nodes
 - ◆ TRIUMF
 - ◇ dCache upgraded to v 2.13.51

Tier 0 News

- There have been issues with the VOMS service regarding notifications for re-signing the AUP
 - ◆ A bug has been identified and will be fixed
 - ◆ Ad-hoc corrections have been applied for now
- In the next days there will be 20k additional cores deployed in the HTCondor cluster

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- The activity levels typically have been high or very high
 - ◆ New record of 111k concurrent jobs reached on Feb 15
- No major issues on the grid
- A new version of AliEn was put into production
 - ◆ In particular it checks and logs the CVMFS status on the WN
 - ◇ To ensure that user jobs have the required SW available

ATLAS

- The new simulation production campaign MC16a started. Reprocessing of ~7 PB of data is about to start, 100k slots are asked for it for about 2 months. Production for Upgrade studies is also ongoing with high memory demands, >5GB memory/core, only a few site can handle. We are working on getting HPCs to dedicate some slots to help with this busy period.
- As one of the mitigation to disk space shortage we are evaluating using the TAPE for task input, namely running derivations from AOD input on tape. The current throughput for reading from all tape sites is about 2.5GB/s, or 200TB/day, which seems tight for order of 5PB of input to process (1month). Improving the tape throughput is definitely possible but very difficult and it will take time (maybe ~6 months).
 - ◆ *Alessandro:*
 - ◇ *results will continue being shared with everyone*
 - ◇ *also other experiments profit from tape system improvements at shared sites*
 - ◇ *optimizing the queue lengths per site is not easy*
 - ◇ *we can do such exercises together with CMS at some point*
 - ◆ *Christoph: CMS had no time for it yet; we would try staging for real production*
 - ◆ *Alessandro: we can follow up on the fts-steering list, convenient for parties involved*

CMS

- finished DIGI-RECO campaign for Moriond17 (total volume 10B evts)
- completed re-miniAOD of 2016 data during February
- Phase 1 and 2 upgrade Monte Carlo generations in progress
- analysis backlog worked down thanks to lower production activity
- we are preparing to move the central services of the CMS global pool / workflow management system back to CERN
 - ◆ new HTCondor version is less resources demanding
 - ◆ high-performance VMs provided by CERN still under evaluation
- EOS issues: file metadata-inode association lost during nameserver failover T0_CH_CERN GGUS:126358 and metadata pointing to wrong inode T2_CH_CERN.
- moving forward on IPv6 and CentOS 7
 - ◆ we are now technically ready to run Release Validation on native CentOS 7
 - ◆ *Christoph and Stephan:*
 - ◇ *currently CentOS 7 can be used together with Singularity to provide an SL6 environment*
 - ◇ *the online teams are moving things to CentOS 7 at P5, which helps for the momentum*

- Had recently some sites that run into issues with check sums used by FTS3 transfers
 - ◆ Seems to be an issue of too Globus GridFTP server not dealing properly with requested adter32 sum
 - ◆ Some relation to recent FTS3 updates/re-configurations at CERN?
 - ◆ *Andrea M: to be followed up on the fts-support list*

LHCb

- High activity levels (from 70K to 90K jobs in average)
- VOMS was suspending user ship for most of our users because of a Acceptable Use Policies expiring on Tuesday. It was a bug in VOMS. An Alarm ticket was raised and all the AUP signatures have been restored for all users which expired at the same day.
- The Oracle DB migration and security patch went smoothly without any problems or significant consequences.

Discussion

Ongoing Task Forces and Working Groups

Accounting TF

- Latest meeting [☞](#) has been held on the 9th of February. Main topic discussed was a possibility to integrate accounting information for the opportunistic resources into APEL. LHCb is quite advanced in this respect. ATLAS and CMS might use a different approach (importing smry data from their experiment-specific accounting systems). However, there are still issues to be resolved in order to make more progress. The main one is benchmarking. Another one is topology description for the opportunistic resources which might be digested by the EGI accounting portal from CRIC. One common problem to be addressed is how to avoid double counting in case info will come both from the site and experiment-specific system. ALICE does not look to be interested in having opportunistic usage accounted by APEL.
- The main topic of the meeting next Thursday is a review of possible implications for accounting in case DB12 benchmark is introduced.
- In parallel started WLCG space storage accounting implementation discussion with the representatives of the DM steering group and WLCG monitoring team

Information System Evolution TF

- Latest meeting [☞](#) has been held on the 23 of February. Agreed with the EGI and OSG colleagues on the policy for introduction of the new service types in the WLCG IS chain in order to ensure naming consistency. Discussed the proposal for the storage service description structure in CRIC.

IPv6 Validation and Deployment TF

- NTR

Machine/Job Features TF

- See the talk in the Indico agenda.

Monitoring

- NTR. Sorry for not being able to attend the meeting, please contact A.Aimar for any monitoring-related matter.

MW Readiness WG

This is the status of jira ticket updates [↗](#) since the last Ops Coord of 20170126:

- MWREADY-142 [↗](#) FTS 3.5.8 for ATLAS & CMS at CERN - completed
- MWREADY-143 [↗](#) FTS 3.6.0 for ATLAS, CMS and LHCb at CERN - ongoing, also LHCb is performing the verification. LHCb discovered a backward incompatibility issue between the previous version of the FTS client and the new server.(Fixed)
- MWREADY-140 [↗](#) ARC-CE 5.2.2 on C7 for CMS at Brunel - ongoing, 5.2.1 verification stopped cause a blocking bug was discovered by devs
- MWREADY-135 [↗](#) WN for C7/SL7 at TRIUMF for ATLAS - on-going, some discussion with CREAM-CE/LB are needed. TRIUMF found that CREAM-CE jobs requires LB client libs installed on the WN, but LB clients are not supported on C7.
- MWREADY-128 [↗](#) UI for C7/SL7 at TRIUMF for ATLAS - on-going, upgrade to FTS 3.5 is needed because of a broken deps. EGI has been contacted.
- MWREADY-141 [↗](#) dCache 3.0.4 at PIC for CMS - on-going. Testing also the new OpenID connect interface. Test with CEPH postponed.

Network and Transfer Metrics WG

Squid Monitoring and HTTP Proxy Discovery TFs

- <http://grid-wpad/wpad.dat> [↗](#) is in production at CERN. It supports IPv6 addresses, but we can't yet enable IPv6 on the servers because although most of CMS has switched to using it, some use cases are still using an old PAC file on the same servers that only supported IPv4. The largest remaining case is expected to be migrated by 9 March.
- <http://wlcg-wpad.cern.ch/wpad.dat> does not yet support IPv6, that is planned to be added later this year.

Traceability and Isolation WG

Last meeting on 2017/03/01 (<https://indico.cern.ch/event/610915/> [↗](#)):

- OSG has made significant progress on testing/integrating/using Singularity:
 - ◆ Singularity deployed in 15 OSG sites, used in more than 1M jobs this week
 - ◆ CMS integration to follow, solution for RHEL7 worker nodes
- Early discussion started on user data workflow for VO

Alessandro: next week there will be an ATLAS workshop on Singularity

Theme: SLAs and usage of different kinds of computing resources

See the presentation [↗](#)

Discussion

- Julia: how can we mark such resources?
- Gavin: their CE could be labeled with corresponding properties
- Andrew: if other sites want to do this, we need coordination and a way to credit the resources
- Alessandro: might sites already be happy filling the holes in their resource usage?
- Maarten: mind that such new paradigms may require a lot of work by the sites; the effort needs to be warranted by some form of credit they care about
- Alessandro: today we have a dependence on resources above the pledge; we can integrate them with proper tags
- Andrew: provisioning above the pledge is recorded too
- Maarten: beware that we cannot simply add the extra walltime provided, as most of the jobs actually might have failed, i.e. no good to the experiment
- Andrew: such resources can be seen as a parallel addition to MoU commitments
- Maarten: better not require that sites do that in a uniform way
- Alessandro: OSG should be involved, e.g. BNL have an extra queue where ATLAS jobs can be pre-empted and which therefore is used for Event Service jobs
- Julia: let's first follow up in a small group to see what can be done at CERN; at the WLCG level we would need to think about how to tag and treat such resources
- **Next steps:**
 - ◆ Follow up in the accounting task force how usage of these resources is integrated in APEL
 - ◆ Come back to this discussion in a couple of months when CERN and experiments accumulate more experience and probably we have other sites doing similar things. Need to address various issues and answer various questions:
 - ◇ Whether/how experiments benefit using these resources? (Benefit of additional power vs effort vs success rate)
 - ◇ Accounting
 - ◇ Generic tagging of such resources so that this information can be used by the workload systems of the experiments

Theme: Machine/Job Features update

See the presentation [↗](#)

Discussion

- Alessandro: the pilot is submitted to an N-hour queue, so it can know the 'N'
- Andrew: a site might advertise a 15-minute queue if an experiment can make intelligent use of such a queue; it could then even be used for pledged resources
- Andrew: LHCb cannot force the sites to provide MJF support
- Julia: a GGUS ticket campaign could be the way to go
- Andrew: when our interruptible MC becomes mainstream, LHCb has a strong argument
- Alessandro:
 - ◆ the MJF framework has nice properties, but ATLAS currently are more concerned with the Event Service approach, i.e. let jobs save results often

- ◆ the pilot asks for cores, wallclock time and memory; ATLAS prefer running the fast benchmark in the pilot, instead of trusting the value at boot time
 - ◆ the pilot can read MJF parameters, but they have not really been used so far
 - ◆ a site might create an opportunistic queue to fill the gaps
 - ◆ how can we get the behavior right on cloud or HPC resources?
- Andrew:
 - ◆ MJF allows the pilot to make the decision when to quit successfully, instead of getting killed by the batch system (i.e. job failure)
 - ◆ a cloud VM could set the MJF parameters at boot time
 - ◆ an HPC resource could have an external web server to supply the parameters
- CMS:
 - ◆ it is OK to ask sites to deploy this
 - ◆ the first usage in CMS would be the benchmark info
 - ◆ CMS jobs can already do a soft drain when a shutdown entry is found
- Maarten: ALICE jobs could make use of MJF functionalities, but that would be a low-priority enhancement, i.e. will not happen soon
- Gavin: at CERN MJF has been removed from LSF, whereas HTCondor will have it soon
 - Andrew: installing the rpm will automatically give sites various functionalities; it is not a problem when some parameters are absent
 - Julia: we should concentrate on LHCb sites first; we will have a WLCG visitor who could help with the deployment campaign
 - Andrew: of the T1 sites IN2P3-CC is missing
 - Catherine (in the Vidyo chat): we will look into what can be done there
- **Next steps:**
 - ◆ An upcoming visitor of the WLCG operations team will be asked to help Andrew in a deployment campaign concentrating on the LHCb sites. Need to understand how much effort it would require from someone who can dedicate to it full time and how much benefit LHCb would get.
 - ◆ Understand the outcome of the benchmarking discussion and what impact the introduction of the fast benchmark would have on the MJF deployment, whether it brings any momentum.

Action list

Creation date	Description	Responsible	Status	Comments
01 Sep 2016	Collect plans from sites to move to EL7	WLCG Operations	Ongoing	The EL7 WN is ready (see MW report of 29.09.2016). ALICE and LHCb can use it. NDGF plan to use EL7 for new HW as of early 2017. Other ATLAS sites e.g. Triumf are working on a container solution that could mask the EL7 env. for the experiments which can't use it. Maria said that GGUS tickets are a clear way to collect the sites' intentions. Alessandra said we shouldn't ask a vague question. Andrea M. said the UI bundle is also making progress. Jan 26 update: this matter is tied to the

				EL7 validation statuses for ATLAS and CMS, which were reported in that meeting. March 2 update: the EMI WN and UI meta packages are planned for UMD 4.5 to be released in May
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
03 Nov 2016	Check status, action items and reporting channels of the Data Management Working Group	WLCG Operations	DONE	
26 Jan 2017	Create long-downtimes proposal v3 and present it to the MB	WLCG Operations	Pending	

Specific actions for experiments

Creation date	Description	Affected VO	Affected TF/WG	Comments	Deadline	Completion
29 Apr 2016	Unify HTCondor CE type name in experiments VOfeeds	all	InfoSys	Proposal to use HTCONDOR-CE.		Mostly DONE

Specific actions for sites

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

AOB

This topic: LCG > WLCGOpsMinutes170302

Topic revision: r20 - 2018-02-28 - 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