

Table of Contents

| | |
|---|----------|
| WLCG Operations Coordination Minutes, March 17th 2016..... | 1 |
| Highlights..... | 1 |
| Agenda..... | 1 |
| Attendance..... | 1 |
| Operations News..... | 1 |
| Middleware News..... | 1 |
| Tier 0 News..... | 2 |
| DB News..... | 3 |
| Tier 1 Feedback..... | 3 |
| Tier 2 Feedback..... | 3 |
| Experiments Reports..... | 3 |
| ALICE..... | 3 |
| ATLAS..... | 3 |
| CMS..... | 3 |
| LHCb..... | 4 |
| Ongoing Task Forces and Working Groups..... | 4 |
| gLExec Deployment TF..... | 4 |
| HTTP Deployment TF..... | 4 |
| Information System Evolution..... | 5 |
| IPv6 Validation and Deployment TF..... | 5 |
| Machine/Job Features TF..... | 5 |
| Middleware Readiness WG..... | 5 |
| Multicore Deployment..... | 6 |
| Network and Transfer Metrics WG..... | 6 |
| RFC proxies..... | 6 |
| Squid Monitoring and HTTP Proxy Discovery TFs..... | 6 |
| Tier0 & Tier1 2016 pledges..... | 6 |
| Action list..... | 7 |
| Specific actions for experiments..... | 7 |
| Specific actions for sites..... | 7 |
| AOB..... | 8 |

WLCG Operations Coordination Minutes, March 17th 2016

Highlights

- Please check the T0 & T1 2016 pledges attached to the agenda[☞].
- The Multicore TF accomplished its mission. Its twiki remains as a documentation source.
- The gLExec TF also completed. Support will continue. Its twiki is up-to-date.

Agenda

- <http://indico.cern.ch/event/504970/>[☞]

Attendance

- local: Maria Alandes (chair), Maria Dimou (minutes), Maarten Litmaath (ALICE), Jerome Belleman (T0), Andrea Manzi (MW Officer), Marian Babik (Network WG), Krystof Borkovec (CERN IT Compute Group), Marc Slater (LHCb), Giuseppe Bagliesi (CMS), Nurcan Ozturk (ATLAS).
- remote: Stephan Lammel (CMS), David Cameron (ATLAS), Alessandra Doria, Michael Ernst, Antonio Maria Perez Yzquierdo, Di Qing, David Mason, Andrew McNab, Hung-Te Lee, Renaud Vernet, Jeremy Coles, Catherine Biscarat, Eygene Ryabinkin, Gareth Smith, Onno Zweers, Oliver Keeble, Thomas Hartmann, Ulf Tigerstedt, Victor Zhiltsov, Vincenzo Spinoso, Xavier Mol, Alessandro Cavalli.
- apologies: Josep Flix (PIC)

Operations News

- Next Ops Coord meetings:
 - ◆ April 7th
 - ◆ April 28th (since May 5th is Ascension day)
 - ◆ June 2nd
 - ◆ July 7th

In principle the meetings are the first Thursday of each month but there need to be exceptions like today and in May. In general, we try to sync with GDB dates in a way that mutual input can be timely, useful and content-full for both groups.

Middleware News

- Useful Links:
 - ◆ [Baseline Versions](#)[☞]
 - ◆ MW Issues
 - ◆ Storage Deployment
- Baselines/News:
 - ◆ The EMI software repository has been replaced by the EGI Preview repo (https://wiki.egi.eu/wiki/Preview_Repository[☞]). It will be updated monthly with software released by the PTs.
- Issues:
 - ◆ Advisory on NSS vulnerability broadcasted by EGI SVG (<https://wiki.egi.eu/wiki/SVG:Advisory-SVG-CVE-2016-1950>[☞]). All running resources

MUST be either patched or software removed by 2016-03-22 00:00 UTC. As there is not a list of Services affected, sites admins are asked to check which services are impacted after the upgrade (needs-restarting command) and restart them or just reboot the host. We have discussed within WLCG ops the need to have those advisories more clear in term of software affected, will check with SVG about this.

- ◆ The DPM team has discovered a critical bug in the new HTTP based drain commands implemented in the dmlite-shell and the lcgdm-dav component available in DPM 1.8.10. Using the new drain commands together with lcgdm-dav version <= 0.16.0 can lead to data loss in some circumstances. Sites have been contacted via broadcast and asked not to use the new drain commands (documented at <https://svnweb.cern.ch/trac/lcgdm/wiki/Dpm/Dev/Dmlite/Shell#Newfunctionality:Drain>) and until the fixed components are released continue to use the old dpm-drain command. One site has been affected in production heavily (200TB data loss)
- ◆ Still issues reported with software using MD5 signature (latest openjdk's have disabled it). This time an old version of the fts2 client (used by CMS) packaged in osg 3.0 with a dependency towards a version of gridsite using MD5 and deployed on some Phedex box was the cause of FTS transfer failures. The affected OSG sites have been asked to move to the version of fts2-client included in OSG 3.2.
- T0 and T1 services
 - ◆ BNL
 - ◇ FTS upgraded to v 3.4.2
 - ◆ CERN
 - ◇ FTS upgraded to v 3.4.2
 - ◆ JINR-T1
 - ◇ dCache upgraded to v 2.10.56.
 - ◇ Postgres upgraded to Postgres v.9.4.6
 - ◆ IN2P3
 - ◇ dCache upgraded to v.2.13.26 on pools
 - ◆ NDGF
 - ◇ dCache upgraded to v 2.15.1
 - ◆ NT-T1
 - ◇ SURFsara will move all grid hardware to a new datacenter in October 2016. We expect to be down for 2 weeks.
 - ◇ Upgrade to dCache 2.13 planned before 2016-05-01.
 - ◆ RAL
 - ◇ FTS upgraded to v 3.4.2

Tier 0 News

- Following CERN's Invitation to Tender IT-4143/IT, a contract has been placed with T-Systems. The idea is to deploy these resources in more similar a way to the standard Tier-0 processes than with previous public cloud procurements. Mass delivery is foreseen for mid May to mid August. We are in contact with the experiments to discuss details.

Maria A. asked if this point affects the pledges or it refers to the Cloud procurement activity, which has no effect on pledges. The answer will be added here when known.

- The loan of CPU capacity to AMS has been extended until March 10, at which time it will be claimed back. There was no impact on the capacity allocated for LHC VOs.
- The HTCondor pool has been extended to 15'000 cores; further extension will require the Kerberos support, which is very close to being available as a result of a very good collaboration between CERN and the HTCondor team. We will invite pilot users and then adapt the ratio of HTCondor to LSF according to user demands.

- The upgrade of computing hypervisors to valid configurations (Kilo-1 wherever possible, EPT switched on on the others) to contribute to better performance has been finished on schedule - thanks to the cloud team!
- The CASTOR software have been consolidated; release 2.1.16 merges CASTOR with the SRM stack, which will allow for co-locating the stager and the SRM daemons. This feature will be rolled out to the new headnodes.

DB News

Tier 1 Feedback

Tier 2 Feedback

Experiments Reports

ALICE

- mostly high activity
 - ◆ also thanks to the HLT clusters: up to 9.2k job slots
 - ◆ new record reached March 10: 99372 running jobs!
- various VOMS-Admin issues caused a support load on the VO admins

ATLAS

- Reprocessing of 2015 data has been completed.
- Large scale MC15c Monte-Carlo reconstruction campaign started, will process 5B events for the next 3-4 months.
- High Level Trigger reprocessing is running.
- Heavy Ion express stream processing is ongoing, bulk processing is expected in ~2 weeks.
- VOMS: VOMS2 has a different behavior wrt VOMS. After one year of deployment, now we have seen a lot of ATLAS users with membership suspension because of the fact that they did not re-sign the AUP. While we know this is to be dealt within ATLAS, it's quite a painful exercise.
- FTS: New release 3.4.x deployed. ATLAS will start consuming messages. Still some duplicates have been noticed, FTS team is aware.

Andrea M.: The FTS issues are not related to the messaging but to a race condition. The problem is known. Tentative fix in the pilot now. Duplicate messages are related to duplicate actual transfers. It is good to consume messages, this will bring the load down.

CMS

- General Status
 - ◆ CMS Detector is being commissioned for 2016 running
 - ◇ Global runs with increasing number of components
 - ◇ First field-off Cosmics taken
- Overall production/processing load is rather low
 - ◆ Ongoing campaigns are in their tails
 - ◆ New major 2016 MC production campaign expected to be launched early April
- Enabling Tier-2 sites for multi-core pilots in progress
 - ◆ 8 sites (19,050 cores) completely transitioned
 - ◆ 6 sites (7,500 cores) work in progress

- ◆ 10 sites (21,600 cores) site admins contacted
- Phedex update: sites are asked to move to Phedex 4.1.5 (required), 4.1.7 is recommended
 - ◆ 1 Tier-1, 5 Tier-2, 22 Tier-3 sites need to upgrade
 - ◆ Sites not at required version received a GGUS ticket
- md5/FTS issue now understood: legacy proxies signed with MD5
- Still fallout triggered by the glibc update/reboot a few weeks ago,
- SAM instability last week/this weekend after VM migration
- SAM test infrastructure will switch from Nagios to ETF by the end of March

Giuseppe added that CMS users also suffered from low voms-admin performance.

LHCb

- Had a clash of 3 Tier 0/1s being in downtime simultaneously earlier this week which would have significantly affected operations.
- Validation Productions for Turbo and 2011 ReStripping currently being processed
- Once validation is complete this evening/tomorrow, major Run1 stripping campaign will be launched early next week and will take ~1 month
- Assuming this goes OK, we will start staging the 2012 data soon after
- In addition, significant MC production and user jobs running.

Marc added the continuing LHCb concern of simultaneous downtimes by more than 2 Tier1s. Maria A. explained the new version of GOCDDB downtime calendar, very helpful for checking clashes. This is linked from the 3pm twiki. The Tier1s were asked to inform the wlcg-scod about downtimes between Monday 3pm calls.

Maarten explained the various issues related to voms-admin, affecting all experiments. One year ago, we were running VOMRS, which had a different behaviour. This is the first year that voms-admin has to handle periodic expiration of AUP signatures and other forms of VO affiliation expiration. The massive co-signature of the AUP by all VO members is not justified to bring the service down. The load shouldn't be that much. There are 2 services, `lcg-voms2.cern.ch` and `voms2.cern.ch`, that should be completely equivalent, but sometimes one of them presented instabilities recently.

Ongoing Task Forces and Working Groups

gLExec Deployment TF

- the TF has been concluded

Maarten confirmed that support will continue. No 100% alternative solution is ready at this moment but significant progress is expected in the course of this year. The GDB recently instated a WG on Traceability and another one on Security Ops Centres which will concern themselves with these matters. The Amsterdam GDB kicked-off these WGs. Ian Collier, GDB chairman, is active in these activities. Experiments are represented, e.g. Alessandro Di Girolamo from ATLAS. Marian asked if gLExec monitoring should be stopped for ALICE, ATLAS and LHCb. Only CMS use it in production. The Ops team will check that with the other experiments.

HTTP Deployment TF

- TF has completed what we hope is the final round of ticketing - ~20 issues still open <http://cern.ch/go/h8Kl>
- Advice on storage configuration has been documented - https://twiki.cern.ch/twiki/bin/view/LCG/HTTPTFSAMProbe#Storage_systems_specific_advice

- ◆ This is particularly important for DPMs
- As most systematic issues have been found and fixed, followup can pass to the experiments when ETF goes live (which will automatically deploy the HTTP TF monitoring).
- Next meeting (Wed 23rd) will discuss what final steps the TF needs to take
 - ◆ <https://indico.cern.ch/event/509854/>
- One proposal is the continuation of the storage provider collaboration beyond the lifetime of the TF, in order to synchronise development targeting mid/long term WLCG data management. Feedback on this proposal is welcome.

Discuss offline, including also Julia, the possible conversion of the TF into a WG.

Information System Evolution

- List of primary information sources is now summarised in this document.
- CMS and ATLAS agreed to evaluate together a common information system (CRIC). First meetings are taking place. It was agreed to work on a prototype in the next few months.
- The strategy to stop depending on the BDII and using GOCDB/OIM as unique information sources will be evaluated as part of the CRIC work.
- Short and medium term plans within the TF will be discussed at the next meeting taking place on 31st of March.

IPv6 Validation and Deployment TF

Machine/Job Features TF

- Finalized specification published as HSF-TN-2016-02 [↗](#)
- Next step is to update and complete implementations, starting with Vac/Vcycle for VMs (done), PBS/Torque (done), and HTCondor (started).
- See MachineJobFeaturesImplementations for links to sources, RPMs, YUM repo etc.
- Will begin rolling out sites, initially with sites that have volunteered to help test updated implementations.
- Update MJF SAM tests.

Middleware Readiness WG

- The 16th MW Readiness WG meeting took place yesterday March 16th 2016. Agenda https://indico.cern.ch/e/MW-Readiness_16 [↗](#), notes MWReadinessMeetingNotes20160316, Summary:
- **Tier1s are invited to tell the e-group wlcg-ops-coord-wg-middleware@cern.ch whether they agree to install the pakiti client on their production service nodes, so that the versions of MW run at the site be known to authorised DNs site managers taken from GOCDB and expert operations' supporters.**
- SRM-less DPM test on-hold until ATLAS pilot code is changed as per JIRA:MWR-104 [↗](#)
- Excellent progress with gfal2 testing (various configurations) as per JIRA:MWR-101 [↗](#) & JIRA:MWR-117 [↗](#)
- Proposed date for the next meeting is **Wed May 18th at 4pm CEST.**

Multicore Deployment

Antonio Y. confirmed, on behalf of the TF that it can now be closed. Advice to LHCb and T2s will continue. The twiki will be updated so that it stands as documentation for the future.

Network and Transfer Metrics WG

- ICFA SCIC meeting was held at J-Park in February, slides from the report (including WG contribution) can be found at <http://icfa-scic.web.cern.ch/ICFA-SCIC/meetings.html>
- LHCOPN/LHCONE Meeting held in Taipei (<https://indico.cern.ch/event/461511/>)
- WLCG Network Throughput SU: ASGC connectivity
 - ◆ Packet loss and high latency for certain packets (queuing issue ?) reported by perfSONAR on ASGC to CERN, but not confirmed by the counters
 - ◆ Narrowed down to the StartLight to ASGC segment, but unfortunately there are very few sonars in Asia with very limited peering, which will impact further investigation
- Stable operations of the perfSONAR pipeline (collector, datastore, publisher and dashboard)
- Throughput meetings held on Feb 24th and March 9th :
 - ◆ Soichi Hayashi presented the new configuration interface that will become part of perfSONAR 3.6
 - ◆ Shawn presented the way we currently monitor the perfSONAR infrastructure, including OSG production services
- perfSONAR 3.5.1 released, 184 instances were auto-updated, only 13 instances on 3.4

RFC proxies

- NTR

Squid Monitoring and HTTP Proxy Discovery TFs

- lhchomeproxy.cern.ch is now being used for LHC@home cvmfs, and monitoring is set up. For now the URL is hard-coded but by next month it should be using Web Proxy Auto Discovery so it can switch to other worldwide proxies, automatically ordered by geo location. Squid deployments are being pursued at Fermilab and in China, Australia, and Brazil.

Tier0 & Tier1 2016 pledges

All slides linked from the agenda. Clarifications from the discussion at the meeting, where needed, are:

- CERN: The resources pledged till April are respected. New capacity is arriving in May. These are opportunistic resources, not related to the pledged figures (? - to be clarified). Not clear right now if all resources pledged for 2016 will be there on time. Maria A. asked to clarify the relationship between the statement in T0 report about massive delivery in May-August and the second statement in the pledge slide.
- ASGC: Oliver said that the DPM performance limitations must be due to some non-optimal configuration because other sites can work without bottleneck issues. To be taken offline. Alessandro says that network issues can harm the performance that achieved pledges can offer. Still dealing with hardware equipment issues related to their network set-up with the vendor. Action on the site to report progress in this meeting.
- BNL: ESnet offers a very good quality network connectivity. Opportunistically, the network capacity can go up to 100Gbps on LHCONE and OPN

- CNAF: Resources available as pledged. **There will be slide uploaded tomorrow, explaining the numbers.**
- FNAL: Resources available as pledged.
- GridPP: Not applicable.
- IN2P3: Resources available as pledged.
- JINR: CMS-only site. Expected increase of about 50%. All resources are hoped to conform to pledges, tape will take longer to increase capacity.
- KISTI: not connected
- KIT: Alessandro asked about the tape infrastructure because slow performance is observed at times. Thomas replies that there is no specific plan for tape system migration from TSM to HPSS, probably it will happen towards the end of 2016.
- NDGF: More memory installed, more jobs can be run, more storage available, hardware was delivered to various NDGF sites at the end of last year and early this year. Additions were transparent. No more expected storage deliveries this year. Tapes are not evenly distributed, so some capacity limitations can be temporarily faced while the overall backup capacity is still abundant. No problem envisaged to fulfill the pledge. Maria A. reminded that monthly Accounting reports gives the current usage. This can be compared with the pledge.
- NL-T1: The site support ALICE, ATLAS and LHCb. The slides list pledges per VO. **The experiments are invited to inform NL_T1 how they wish to distribute the increase of storage between disk and tape.** April pledges are fine with some exception concerning ATLAS pledged disk. There will be a downtime in April to replace dCache database servers, and a long shutdown of 2 weeks in October.
- NRC-KI: All on slide. No comment.
- OSG: Not applicable.
- PIC: All on slide. The T10KD tape incident is not yet understood.
- RAL: All on slide. Resources available as pledged.
- TRIUMF: All on slide. Resources available as pledged.

Action list

| Creation date | Description | Responsible | Status | Comments |
|---------------|---|-------------|--------|----------|
| 17.03.2016 | The Operations' team to check with ALICE, ATLAS and LHCb whether the gLExec monitoring can be stopped. | Maarten | New | |

Specific actions for experiments

| Creation date | Description | Affected VO | Affected TF | Comments | Deadline | Completion |
|---------------|-------------|-------------|-------------|----------|----------|------------|
|---------------|-------------|-------------|-------------|----------|----------|------------|

Specific actions for sites

| Creation date | Description | Affected VO | Affected TF | Comments | Deadline | Completion |
|---------------|---|-------------|-----------------|--|----------|------------|
| 17.03.2016 | The Operations' team to follow-up ASGC's progress on the networking issues experienced there and report progress at the Ops Coord meeting | ATLAS | network metrics | | | |
| 22.01.2016 | CMS sites are requested to move to Phedex 4.1.5 (minimum) or 4.1.7 (recommended) on SL6 | CMS | - | Should be done by the end of Feb. GGUS | - | CLOSED |

| | | | | | | |
|------------|--|-------|---|--|---|--------|
| | | | | tickets are opened. Still 9 sites missing. | | |
| 18.02.2016 | CVMFS monitoring: Taiwan Stratum 1 is always in trouble, better to turn it off than have it out of date? GGUS:119557 ↗ | ATLAS | - | Done | - | CLOSED |

AOB

-- MariaDimou - 2016-03-15

This topic: LCG > WLCGOpsMinutes160317

Topic revision: r62 - 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