

Report 18/02/2016

- See presentation in agenda [↗](#)

Report 07/01/2016

- Accounting:
 - ◆ John Gordon update: The default EGI has not changed but WLCG users should mainly use the Tier1 and Tier2 views (eg <http://accounting.egi.eu/tier1.php>) which now use the same data as the production portal (ie include cores). The EMI3(WLCG) view also includes cores and would be useful to view an integrated view of a country including both its Tier1, Tier2, Tier3 and other sites.
 - ◆ On ATLAS side working on comparing accounting records in the dashboard and in APEL site by site for the T1 and region by region for T2s.

Report 16/07/2015

- Accounting: Latest update on accounting sites that haven't enabled yet multicore accounting was sent to the GDB mailing list on the 6/7/2015 with the following list of sites
 - ◆ Austria - HEPHY-UIBK, Hephy-Vienna
 - ◆ Germany - RWTH-AACHEN, DESY-HH, MPPMU
 - ◆ India - IN-DAE-VECC-02
 - ◆ Mexico - ICN-UNAM
 - ◆ Russia - RRC-KI, ru-Moscow-SINP-LCG2, RU-SPbSU, Ru-Troitsk-INR-LCG2
 - ◆ Spain - IFIC-LCG2, UB-LCG2
 - ◆ ~~UK - UKI-LT2-IC-HEP, UKI-SCOTGRID-DURHAM, UKI-SOUTHGRID-OX-HEP~~

could the sites listed check and let us know (UK sites we know about that is why I stroke them off)?

Report 02/07/2015

- Accounting: John Gordon will send out a broadcast and produce an updated list of missing CEs. Action should stay until all sites done.

Report 18/06/2015

- Multicore accounting: several sites haven't yet enabled on some or all their CE yet. mostly CREAM-CEs but some ARC-CEs appear here and there too. APEL team has opened tickets for the NGIs. But here is a reminder of what the WLCG sites should do. List of computing elements. If any problem contact the APEL team.

Report 04/06/2015

- ATLAS deployment:
 - ◆ Reminder that the goal is to have 80% of production resources usable by multicore and sites should modify the configuration to respect that. See Actions for sites.

Report 21/05/2015

- ATLAS deployment:

- ◆ Goal: 80% of production resources usable by multicore. This corresponds roughly to 150k slots. The max achieved has been 110k slots. It is not a problem of sites not having a queue as only 9 smaller sites are still missing. It's a matter of tuning the system.
- ◆ ATLAS is looking at increasing the jobs length up to about 10-15h for 8-core slots to improve on the efficiency and is also looking at a global fairshare to avoid sending too many low priority single core jobs when there are high priority multicore in the system.
- ◆ Sites should instead look at their setup and try to prioritize multicore over single core. The base TF solutions can be found here. In particular Torque, Htcondor and SGE have all a solution.
 - ◇ Sites that are already dynamically configured but still cap the multicore slots should also revise this number to respect the 80% share. As a reminder the shares for ATLAS jobs are as follows T1: 5% analysis and 95% production; T2: 50% analysis and 50% production. So multicore should get 80% of 95% at T1s and 80% of 50% at T2.
- CMS status:
 - ◆ Consolidating deployment to T1s: working together with CNAF and CCIN2P3 to increase multicore resources allocation results
 - ◆ Analysis of multicore pilots internal (in)efficiencies ongoing

Report 19/03/2015

Passing parameters to the batch system: last update about this was in November. Since then Atlas added extra memory parameters to pass to the batch systems according to the plan presented at the last S&C week [☞](#). While parameters were already passed to ARC-CE sites, they were with a different memory scheme. The new scheme is now being tested at 3 ARC-CE/Htcondor sites in the UK on their multicore queues. For Cream sites which were the most controversial and debated it is in test at Nikhef and Manchester which both have a cream/torque combination. This has helped ironing out the most macroscopic problems such as units used. In Manchester the new scheme has been now enabled on all the queues on 1 of the clusters, i.e. analysis jobs as well as production multicore and single core and will see in the next few days how it goes. If no adverse effect is observed the plan is to extend to more queues on the sites in test already and contact other sites to test other batch systems/CE combinations.

ATLAS only 13 sites are missing a multicore queue: 3 are too small and 10 have a plan for the next few weeks months.

Report 05/03/2015

The first objective of the TF, that is, the understanding of the principles for a successful shared use of the common resources in multicore mode by ATLAS and CMS has been achieved. The initial deployment to the involved sites (T1s for CMS, T1s+T2s for ATLAS) has also been successful and the most popular batch systems' capabilities concerning the use of multicore jobs have been discussed as well.

However, at this point both experiments are working independently on their respective infrastructures. We therefore propose to keep the TF open in "passive mode" while this is ongoing, in order to review the status once both experiments have advanced in their respective models, and in case common matters need to be discussed during this period.

Report 22/01/2015

- CMS multicore at T1s, see notes above. Deployment to T2s to restart once the submission infrastructure (pilot factory) testbed is deployed.
- ATLAS 26 T2 to enable followed in JIRA (see ATLAS report)

Report 04/12/2014

- CMS multicore:
 - ◆ Ongoing test of submission of PromptReco multithreaded jobs to T1s.
 - ◆ Working with CNAF to understand and improve low number of max. multicore pilots which get to run.
 - ◆ Test deployment to CMS T2s still waiting for test-bed infrastructure (pilot factory) deployment.

Report 20/11/2014

- Passing parameters to the batch system: GDB report [?](#).
 - ◆ Currently completing the table of parameters from the batch systems and testing the CREAM-CE capabilities. The final summary of this and whatever we decide will be in this page.
 - ◇ FZK tested SGE scripts with what is in the current script and it works.
 - ◇ Manchester tested sending random strings to torque: they do get accepted with direct job submission, so potentially we could use Glue2 for CREAM-CE if we really want. Tested also with different operators >= and == and blah does append a _Min suffix and nothing in the second case. So we can restrict just to use == as it should provided the scripts are adapted.
 - Adapting the scripts and how to distribute them hasn't been discussed yet.
 - ◇ CERN script for LSF is heavily used but adopted a different method grouping all the requests in one LSFResource parameter with no reference to any Glue schema.
 - ◇ ARC-CE using RSL [?](#) directly no Glue
 - ◇ HTcondor CE not developed yet. The developers joined the TF this week.
- Accounting: WLCG MB report [?](#)
 - ◆ EGI is now looking into accelerating the progress of the new development portal
 - ◆ To have the correct accounting sites also should:
 - ◇ EMI-3 CREAMs have to enable multicore support
 - Please check 18/09/2014 TF report for instructions
 - ◇ Sites using SSM1.2 should move to SSM2
 - ◇ Sites using DGAS should move to use the APEL client

Report 06/11/2014

- Passing parameters to the batch system: discussion started within the TF, so far only few sites participated.
- ATLAS:
 - ◆ 40% of the production resources used occupied by multicore since the beginning of September.
 - ◆ Still 37 sites to move. Work needed both on sites and atlas side to setup the queues
 - ◆ Atlas very interested in the outcome of the passing parameters discussion. Setting them up on ATLAS side will not take long once parameters decided
- CMS:
 - ◆ Multithreaded reconstruction application ready, first test workflows run at PIC, now moving to submission tests at T1s scale
 - ◆ Deployment to T2s: multicore pilot submission tests to be started with some candidate sites
- Multicore TF presentation next week at the GDB

Report 16/10/2014

- CMS activities:
 - ◆ Testing submission of multithreaded jobs started. Jobs running at PIC in multicore pilots along with single core jobs
 - ◆ Working on improving multicore pilots performance monitoring
 - ◆ Started testing submission of multicore pilots to CMS T2s (running regularly at T1s)
- CHEP15 contribution: plans to submit an abstract to CHEP15 concerning TF activities. Draft currently being prepared by coordinators, to be circulated to TF members.

Report 02/10/2014

- ATLAS: see ATLAS report

Report 18/09/2014

Accounting: Publishing multicore accounting to APEL works. ARC CEs publish correctly. For CREAM CEs to make it work it has to be an EMI-3 CE and it has to be enabled in the configuration.

Edit `/etc/apel/parser.cfg` and set the attribute `parallel=true`.

If the site was running multicore already, before upgrading and/or applying this modification, they need to reparse and republish the corrected accounting.

Report 04/09/2014

- **ATLAS:** 11 T1s and 35 T2s are currently running multicore for ATLAS with an average slot occupancy of 30k slots up to 57k. Among Tier2s production is still dominated by US dedicated sites but extra EU shared sites have been added in the past month. 1 French, 3 UK and 4 Italian Tier2s are among the sites added over the summer. Most of this sites have a dynamic or a mixture of dynamic and a small number of initial reserved slots. All sites are requested to enable dynamic batch system support to run multicore and single core jobs concurrently.
- **CMS:** Running multicore pilots regularly at all CMS T1s. Multicore pilots testing at several US T2s as well.
- **Passing parameters to the batch systems:** as reported in other occasions to apply certain scheduling techniques it is necessary for experiments to pass parameters to the batch systems ahead of being scheduled. To do this there have been many obstacles which need still to be cleared. One of these is the fact that there is not a standard blah script for sites deploying CREAM CE and those that exist are clearly not standardized since were written by sites trying to match what the experiments were sending but without any previous discussion. This has been discussed with ATLAS since there are other use cases for which this is needed for example requesting a specific amount of memory. It was decided the TF would take on board the standardization of the blah scripts (and other CEs scripts if needed) for the scheduling parameters. A discussion about distribution and deployment will also be needed.

Report 24/07/2014

- ATLAS: restarted multicore production and up to now has had a pretty stable flow of jobs [↗](#) for the past 2 or 3 weeks.
- CMS and ATLAS have been running concurrently at 5 T1 sites. Here is the monitoring for ATLAS PIC, FZK, IN2P3, RAL and CNAF [↗](#) and CMS FZK [↗](#), RAL [↗](#), IN2P3 [↗](#), PIC [↗](#), for this period.

- ◆ PIC farm, for example, has been consistently running CMS, ATLAS T1 and ATLAS T2 multicore jobs for three weeks. The principle of dynamic allocation of WNs, provided here by the mcfloat script, is showing good results, as jobs are running together reusing empty multicore slots, while the total number of WNs allocated to multicore jobs has been slowly but steadily increasing along the weeks with good overall farm utilization.
- UK: multicore discussed in depth in UK this week. SGE, HTcondor and Torque are the most used batch systems and there is a solution to try for each. A site for each batch system currently running ATLAS jobs. The other sites asked to have links to different solutions and more details on the configuration from the TF page which is on the TODO list - in particular the solution for Torque (mcfloat script as used now by NIKHEF and PIC) need more detailed explaining. Italian sites asked for the same at the WLCG workshop.

Report 07/07/2014

Report at the WLCG Workshop in Barcelona [↗](#)

Report 19/06/2014

- CMS has been running multicore pilots for single core production jobs at most CMS T1s (PIC, KIT, RAL, JINR and CCIN2P3) for over a month now. Since we are sending a mix of single core and multicore pilots to pull jobs from the same pool, whenever there is workload for the site (which is most of the time recently), multicore pilots are sent. So this means actually that the sites are constantly receiving a relatively stable amount of multicore pilots from CMS.
- With respect to feedback from the sites regarding this activity, we don't have a detailed report from any of them yet, but no complaints either so far. We are expecting to collect reports from a few sites (for example KIT, RAL and PIC) in a coming Multicore TF meeting in order to be ready for the coming WLCG workshop.
- ATLAS hasn't restarted multicore submission yet due to problematic software validation.

-- AlessandraForti - 19 Jun 2014

This topic: LCG > MulticoreTFReports

Topic revision: r37 - 2016-02-18 - AlessandraForti



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