



Accounting Review

Summary and action list from the (pre)GDB

Julia Andreeva CERN-IT
WLCG MB
19th April 2016

The goals of the accounting review

- We have quite some issues and open questions related to the accounting
- Effort is dispersed across various task forces and working groups (IS task force, benchmarking WG, cloud procurement work, WLCG resource reporting task force, etc...)
- Need to summarize issues we have, define an action plan and propose how we can better organize this work

Topics which were discussed

- CPU (wallclock) accounting
 - Multicore accounting
 - Accounting of the resources provided through cloud interfaces
 - Accounting of the opportunistic resources
 - Review of the EGI accounting portal and accounting reports generated by REBUS
- Storage space accounting. In particular what are our plans for no-srm storage

Conclusions and action list

Multicore accounting

- Seems to be in good shape
- Proper reporting of the number of cores is enabled at all T1 and most of T2. Corner cases to be followed up
- Main effort required for validation of the multicore information in the accounting portal
- Information can be crosschecked with CMS and ATLAS Dashboard where multicore accounting is available

Accounting of the resources provided through cloud interfaces (private clouds)

- The good start was the work done by the WLCG Resource Reporting Task Force. However, it looks like there is still a lot of effort required before we have production solutions for instrumenting our cloud resources for proper reporting
- With OpenStack we progress well thanks to work of the CERN IT-CM group. What about other platforms for cloud computing?
- Need to understand whether other WLCG sites apart from CERN work on it. If yes, need to streamline this effort. And collaborate with EGI.

Accounting of the cloud resources provided by commercial clouds

- Work on accounting of the commercial clouds is in progress and is driven by the CERN IT CM group and in the scope of CERN cloud procurement activity.
- Currently no actions are required.

Accounting of the opportunistic resources

- Some of the experiments believe that it would be nice to have a capability to expose usage of the opportunistic resources in the accounting portal
- These resources are already accounted in the experiment-specific systems which can expose accounting information through APIs
- The proposal is to evaluate a possibility to inject information of usage of the opportunistic resources in APEL, the same way it is done for OSG accounting info which is injected from GRATIA.

Review of the accounting portal and accounting reports

- **Agreed** that though available capacity is required for operations, it does not need to be a part of the accounting reports
- **Agreed** that manual fixes of the accounting numbers for the accounting reports should be avoided and if accounting data is not properly reported to APEL the problem should be fixed there instead
- **Agreed** that WLCG needs a dedicated simplified instance of the accounting UI, which provides customized menus and contains ONLY metrics required by WLCG
- **Agreed** that we use GGUS for reporting bugs and submitting requests to APEL and accounting portal developers
- **Agreed** that better coordination between APEL/accounting portal developers and WLCG represented by the WLCG operations team is required, in particular in what concerns validation of the accounting portal (in terms of functionality) and checking of the trustworthiness of the provided information.

Some more proposed and agreed changes on the accounting reports and accounting portal

- In the reports - use a consistent metric with the one in which pledges are expressed wallclock HS06 instead of CPU HS06.
- In the reports - use the same unit HEPSPEC06 days both for T1 and T2 reports.
- On the portal - add a simple metric 'Elapsed time multiplied by number of cores' in order to easily crosscheck information with the experiment specific accounting systems
- On the portal - mandatory views for the WLCG resources T1,T2, (for all of the formal WLCG sites) and a Country one that includes all sites contributing time to any LHC experiment

Some highlights regarding REBUS

- REBUS is used for pledges, installed (available) capacities and generation of the accounting reports
- Accounting reports are constructed taking data from REBUS itself and accounting data from the accounting portal
- The functionality of REBUS might be reviewed if the new WLCG configuration system is accepted. REBUS can be turned into a service for report generation taking data via APIs from various external sources

Proposal for setting up the CPU accounting Task Force

The goals are :

- to agree with all stakeholders on the content of the accounting portal and accounting reports (metrics, units, aggregation in terms of time and set of attributes like site, VO, country, etc...);
- coordinate with APEL and accounting portal development team and follow up on the progress
- validate the new version of the portal, ideally a dedicated customized UI for WLCG, and make sure that the requirements of the stakeholders are addressed
- validate accounting data, ensure that it is correct in the APEL repository so that manual fixes could be avoided
- follow up on whether/ how we enable injection of the opportunistic resources into APEL and accounting portal
- In case CRIC is accepted as a WLCG central configuration system, re-evaluate role and functionality of REBUS which can be turned into a simpler service responsible for pledges and report generation
- assess the situation with accounting of the pledged resources provided through private clouds , understand whether we have all components in place to properly account such resources , make sure that they are correctly publishing to APEL. In case this requires substantial development, testing and evaluation effort , this can be a question of launching its own dedicated task force

Storage space accounting

- There are two main use cases for storage space accounting which experiments currently deal with:
 - Very detailed space accounting with a possibility of the full storage dump. Implies experiment-specific requirements and data organization
 - High level space accounting (few numbers : total/used/free per VO and where/if possible with some level of details for high level quotas)
- As much as WLCG infrastructure is concerned the second objective is more important and looks to be easier to address in a common way

Steps towards generic WLCG storage space accounting

- Agree with the experiments on high level accounting requirements (second use case on the previous slide)
- Agree with the experiments on the common format of the 'space usage record' (second use case on the previous slide)
- Inspect solutions we have for various storage implementations in order to generate 'space usage records'. Whether we have solutions for all storage implementations which are currently in use? Whether these implementations are optimal?
- Work with storage providers in order to enable generation of the required space usage records
- Understand how we collect, store and visualize storage space accounting information, possible alternatives:
 - Work with EGI in order to enable this info in APEL and accounting portal
 - Investigate WLCG-specific solution, possibly based on ATLAS space accounting Dashboard

How we organize work on storage space accounting

- This work was started by the WLCG operations. Discussions were held with ATLAS, CMS and LHCb
- Further effort can be organized in a form of a dedicated task force or working group
- Another possibility would be to handle this work under the umbrella of the WLCG storage providers working group, which is needed not just for this task

In whatever form this work is organized, the goal is to assess the possibility to have a WLCG solution for high level space accounting which is not based on SRM

- evaluate existing solutions in the experiments
- agree on common formats for space usage records
- collaborate with storage providers and experts in the experiments at sites in order to enable usage record generation scripts for various storage implementation and make sure that there are shared in the community
- collaborate with APEL development team to understand data collection and exposure through the EGI accounting
- assess other scenarios for data collection and aggregation (WLCG-specific repository)
- based on this assessment come up with the proposed architecture and estimation of the required development effort