



CERN Accounting

Gavin McCance
27 September 2018



Accounting scope

- We're collecting accounting from:
 - HTCondor batch, 35 schedds over 2 pools
 - Legacy LSF batch, 2 instances
 - SLURM HPC, 3 instances
 - BOINC native and BOINC/HTCondor
 - Cloud services and compute
 - IT services to experiments
 - e.g. ElasticSearch, Ixplus, databases, dashboards, ...
-

Accounting scope

- We collect wallclock, CPU and (soon) storage
- For most things we charge utilised wallclock. For "dedicated" things, we charge quota wallclock (regardless of utilisation)
- Reports made for CERN departmental accounting and for WLCG (APEL)
 - WLCG is ~60% of total
 - Some item charged to pledge, some items reported "above the line"

Accounting status

- We've spent quite some time reworking the CERN accounting to absorb the new use-cases (we only used to report LSF)
- Still some issues and lessons learned
 - Mostly due to data handling issue and data quality...

Completeness

- Incomplete data is a problem for accounting
 - Down and deleted hosts, late data, mishandled data
 - Mostly now taken care of in our procedures
 - Alerts to highlight persistently missing data
- Currently our WLCG numbers are suffering from an early assumption in our data collection that finished HTCondor jobs appear in a timely fashion in its accounting files
 - ..in fact, many jobs appear up to 10 days later
 - They weren't counted because we assumed that yesterday's file continued all of yesterday's finished jobs

Data validation

- All data items should be sanity-checked before adding them to a monthly sum :-)
- Currently our WLCG numbers are suffering from a few thousand short docker jobs in HTCondor, returning an int64 CPU time from an uninitialised struct
 - ...some really big +ve and -ve numbers result in exciting efficiencies
 - Bug will be fixed upstream but we should have been rejecting stupid records

Issue status

- Late data
 - the new code is much more resilient against late data arriving - new data triggers a recalculation of that days/months' accounting
- Bad data:
 - sanity checks are being put in place for data records
- The fixes should go into production soon

Auto-validation

- We discovered some problems by comparing to ATLAS Panda data
- The fixed code agrees now with the ATLAS data
- Challenges in comparison
 - Empty pilots
 - Timezones(!)
 - Actual node normalisation vs. mean
- Working to run this comparison automatically every day and highlight discrepancies early