

# CERN Batch Accounting

Status update on verification efforts

Gavin McCance, Gorka Bravo Martinez, Florentia Protopsalti, Jaroslava Schovancová  
CERN IT

# Previous talks

- WLCG Accounting TF, Sep 2018: <https://indico.cern.ch/event/758334/>
- WLCG Accounting TF, Mar 2018: <https://indico.cern.ch/event/711469/>

# Accounting scope

- Collecting accounting data 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, ...
- We collect **wallclock**, **CPU** and (soon) **storage**
  - Mostly: charge utilised wallclock. Dedicated resources: charge quota wallclock
- 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

- Spent some time refurbishing the CERN accounting to absorb the new use-cases (we only used to report LSF)
- Experienced some issues and lessons learned
  - Mostly due to *data handling* issue and *data quality*...

# Completeness

- Incomplete data - a notable 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
- **Fixed:**  
our WLCG numbers were 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 appeared 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

- Our WLCG numbers were 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 fixed in the upstream & deployed
  - In addition, we are rejecting suspicious records

# Current Status

- Late data
  - new much more resilient code deployed to production in Sep 2018
- Bad data
  - sanity checks implemented, CI/automation in progress

# Automated checks

- Discovered issues while comparing with ATLAS PanDA data ~1 yr ago
- The fixed code (in production) agrees with ATLAS data
- Challenges to compare
  - Empty pilot wrappers
  - Timezones handling
  - Node normalization vs. mean
- In progress: automation of the comparison
  - to have daily checks
  - to trigger alert early when discrepancies found
- Set up a “console dashboard” to visualize and spot issues





# CERN Batch Accounting

Status update on verification efforts

- Experienced issues in the past, now fixed.
- Automated checks to spot new issues early in progress.

Gavin McCance, Gorka Bravo Martinez,  
Florentia Protopsalti, Jaroslava Schovancová  
CERN IT

