Resource pledging process
Comments from ATLAS

Torre Wenaus, Davide Costanzo

WLCG MB meeting
June 19 2018

Comments

- No objections to CMS proposals
- Regarding granularity of pledges, main issue for us is accurate information on what resources have been provided
  - We've lately been assessing what our T1s are actually providing and it is non-trivial to get the numbers
  - Cannot rely on installed capacity numbers in Rebus
  - Basically we depend on the goodwill of sites/FAs to tell us what we'll get when, when pledges aren't there on April 1
  - We need a better way!
- Ability to declare quarterly pledges could have benefits, even though our usage is flat
  - e.g. if a Tier-1 is late in coming up to pledge, a FA could balance out a low level in Q1 with a correspondingly high level in later quarters
- A coming issue is how do we best cite and measure resource levels -- HS06 is dated and HS17 isn't necessarily the optimal replacement
  - The real metric of interest is events/sec
  - ATLAS is considering developing a realistic, packaged (containerized) test suite with simulation, reconstruction, derivation, evgen standard workloads to measure event throughput as the metric
    - Collaboratively with others?
  - Our requirements are really expressed in event counts; translate to resource requirements via such a metric
  - Not easy but worthwhile, particularly as more non-traditional platforms play a role in fulfilling resource needs