

Supporting US users
(How do we fulfill our WLCG
pledges and make US CMS happy
at the same time?)

Discussion Leader:

Burt Holzman (Fermilab)



The Charge

“Hi Burt,

One thing I want to push on in the coming months [is] making sure that we carve off resources that are in excess of our pledge to WLCG and give US CMS first priority on them. This is of course subtle [...]

Best wishes.

Ken”

[NB: We're discussing Compute resources only in this discussion]





And one more note

This is a discussion!





Fulfilling the pledge

- Pledges to WLCG are reported and monitored on a monthly basis
- I can think of three options that we could implement





Option 1: Static Allocation

- Steady state and straight-forward
- Count up 1500 kSi2k worth of nodes, guarantee fair-share on them
- Prioritize USCMS use on the remainder
- Not entirely straightforward (and not always steady-state): our clusters are heterogeneous





Option 2: Dynamic Allocation

- Start with entire cluster configured for fair-share
- Monitor OSG accounting information daily (or more) and determine when monthly pledge has been fulfilled
- Flip the switch and give US CMS priority on entire cluster
- Reset to fair-share on first day of new month
- Does this meet our users' needs?





Option 3: glideinWMS

- Mandate all jobs be routed through glideinWMS
 - central control, no work needed at sites
- Unfortunately:
 - Not yet enforceable (we run vanilla and gLite WMS jobs)
 - Requires some development – we have the controls to tweak VO-wide priority, but not an application to monitor and adjust it to our needs
 - So this is a non-starter, but I would have to fire myself as glideinWMS lead if I didn't mention it





Prioritizing US CMS

- How do we distinguish between US CMS use and non-US CMS use?
 - VOMS role: /cms/uscms (but users have to remember to use it!)
 - Global glideinWMS analysis CRABserver may cause confusion
 - Do we need a separate US Crabserver/glideinWMS ?
 - Do we need some glideinWMS development to manage separate CMS from US CMS service proxies?





Prioritizing US CMS (2)

- Burt's Accounting group patch in production at FNAL, MIT and submitted to VDT (5838)
- Other accounting group patches enabled at UCSD, Nebraska, Florida , Purdue
- Should we standardize? (even if we weren't doing this work?)





Accounting Groups

- Option 1 (steady-state):
 - Set startd SlotWeight equal to $kSi2k/core$
 - Set GROUP_QUOTA to $(Sum\ kSi2k - 1500)$ for USCMS group, with GROUP_AUTOREGROUP True
- Option 2: (dynamic):
 - When pledge is fulfilled, set GROUP_QUOTA for USCMS to entire size of the cluster





Other questions

- How do we monitor it?
- How flexible are we to changes (pledge changes, new hardware, ...)?
- Is it worth the effort?

