

Proposals for discussion

- Allow some jobs to request memory over physical amount per core
 - up to site configured max RSS
 - unlimited number of cores if mix means no idle cores, or capped at low(10%) level?
 - VO responsible for this. Monitored and enforced by the site - how? Trust but verify?
 - accounting for any idle cores important for pledges
 - important enough to develop memory dimension or rescaling in APEL?
 - HS23 does not scale linearly with HT cores. Accounting to reflect this.
- Move to 16 core as new standard(currently 8) where a VO requests it,
 - if then CMS would want 16 everywhere. ATLAS ok with mix, especially when HT-off.
 - major/all VOs send 16 core jobs as standard, to ease slot re-use
 - must allow 1 core jobs
 - ALICE want then 72hr walltime. CMS also prefer longer walltimes for more cores.
 - to do with draining inefficiency at the end

Whole node scheduling

- Wholenode only where advantageous, e.g. GPU, numa pinning, many-core job(non-parallel)
 - oversubscription of cpu can improve efficiency when some jobs not cpu-bound
 - useful for packing gaps and draining, i.e. better slot efficiency
 - should be allowed also in MCORE if cgroup contained
 - otherwise let the Batch System schedule
- Pledged resources must be able to continue with S/MCORE jobs
 - combining whole nodes with S/MCORE jobs is problematic for many sites
 - If (CMS & ALICE & expert BS admin & volunteer) then do it
 - need at least 2 VOs with reliable whole node jobs, to keep slots
 - for this subset ATLAS would submit some wholenode(but also S/MCORE)

Relaunch Multicore TF to address this between exps, sites and WLCG