

WLCG Memory Requirements

- At the MB at the end of 2015 we agreed that the experiments did not want to require more than 2 GB of RAM per core as a default need, and for the most part will continue to require only 2 GB/core. However, there are clearly several workloads that need or benefit from larger memory machines. There are a number of mechanisms for provisioning such:
 1. Buy machines with more than 2 GB/core. This may be an unnecessary cost if the majority of the workloads still fit within 2 GB. The cost is approximately 20% more to buy 4 GB/core compared to 2 GB. (actually with recent machines with 10 cores per CPU this problem is side-stepped as the proposed configurations have ~3 GB/core)
 2. Switch off hyper threading so that the real cores have twice as much memory available. This is approximately a 20% performance reduction compared to having HT on.
 3. When a job requests more than 2 GB, allocate at execution time 2 cores to provide twice as much RAM, and accept the lower efficiency of the second core not being fully utilised.
- Options 1, 2 require either planning at purchase time, or up-front configuration of the machine. Option 3 is attractive as it allows the site to be flexible in its provisioning from one job to another, and the effective loss of CPU efficiency is acceptable if this is not the majority of the workload.
- All 3 options require a mechanism to request the appropriate resources of the batch system (or equivalent), and to be able to provision correctly.