

The future - Cambridge Perspective

John Hill



UNIVERSITY OF
CAMBRIDGE

- Currently we have
 - 290TB storage which will only decrease as equipment dies. (All except 80TB is already >5 years old).
 - 408 cores CPU (EOL on various dates until early 2021), which will be increased using the GridPP5 money we have already been given
- We are unlikely to get resources for new Tier 2 equipment from outside GridPP.
- The University does provide “in kind” resources at the moment by not charging for electricity, cooling or rack space in the West Cambridge Data Centre (WCDC where the Tier 2 is located). However this may change in the future, and use of the WCDC is not completely free as we have to pay an annual fee for connecting to the network (fixed connection fee + capped usage charge (T2 is 95% of traffic)).
- The University would also like to require all equipment >5 years old to be removed from the WCDC, to encourage the use of more modern, efficient kit. However, there is no “compensation scheme” in place at present to make this approach attractive to users. Cooling and electricity charging might be used to offer an incentive first!

- Effort required to run a “static” small site at a large university is minimal
 - Network is high quality
 - Storage load average typically lower than for a large site and hence fewer problems
 - CPU problems are generally contained and mean at worst a few lost jobs
- Amount of work needed increases significantly when things become “dynamic”
 - Significant new middleware versions need rolling out
 - Moving to a new operating system
 - ...
- The Cambridge group are currently content to provide the required support. **There are also wider benefits from being part of an active HEP sysadmin community.**
- However given the likelihood of little or no capital money for us in future it is difficult to see how this commitment (**and being a good citizen**) can continue beyond the near term, unless something changes.

- One possible area of change might be in the needs of the local users
 - Currently there is (deliberately) no overlap between the local cluster and the Tier 2
 - This simplifies the management, and in any case the spare resources on the Tier 2 are not sufficient to be worth the effort of making them easily available to the local users
 - The users on our two main experiments at present (ATLAS and LHCb) are geared up to using the Grid for most of their work.
 - We are members of MicroBooNE and DUNE. At present our local resources are adequate to support their activities. As the DUNE collaboration develops this may change. However, this will probably result in them using Grid resources in general, rather than requiring a locally enhanced Tier 2.
 - Our theory colleagues mainly use their own private cluster. They have also made limited use of the Cambridge HPC. This pattern is likely to continue in the short term at least. In particular they have shown no interest in using GRID resources.

- So how should the site develop?
 - As said earlier, we are likely to continue to provide support for the Tier 2 for as long as we get capital money. Beyond that it is hard to justify.
 - Remote (to Cambridge) support is a possibility if it is considered desirable for the site to continue.
 - However, this is clearly less convenient (and probably less efficient) than having local support.
 - What about other university resources as a more efficient way to provide a “Tier 2”? There are moves towards a “cloud” environment within the university, but it is probably several years before that becomes a resource we could use routinely. In any case, as with any resources provided by the University of Cambridge, this would cost real money.
 - It may also be possible to make use of the Cambridge HPC. However in the past when this has been looked into, it has been proved to be more trouble than is justified (especially given that we would have to pay for the privilege).

- Conclusions:
 - We are unlikely to provide group resources to support the Tier 2 once we cease getting GridPP capital
 - At the moment there are no other obvious local sources of resources to fill the gaps (both in manpower and to replace the expected run down in GridPP hardware).
 - Remote support is possible for the existing kit, though this is clearly less convenient than local support. GridPP would also need to decide whether they think that the costs which the site would still incur (e.g. networking) justify the continuation of the site.
 - University moves towards a cloud-type provisioning of resources might provide an alternative mechanism for providing the Tier 2 resources. However, it would still almost certainly cost real money and will not be a viable reality for several years.
 - [A batch system change is due](#). Open to deploying VAC or moving to ARC/HTCondor. What can't happen is ending up with both because some VOs support one and some the other. Guidance is needed from GridPP as to whether we should only support ATLAS and LHCb and no one else.