

ATLAS UK Cloud Support meeting minutes, 21 February 2019

Diskless Sites:

We had a discussion about what we (ATLAS UK) want to do about diskless sites, starting from Stephane Jézequel's ADC Weekly [slides](#) (also attached to this Indico page). Here are some of the main points from a lively and lengthy discussion.

- * "Lightweight" sites with small disk can still provide significant CPU. Eg. Durham is listed as having 192 TB (currently 302 TB), but provides 2.7% of ATLAS UK CPU. See [chart](#).
- * We probably have three classes of Tier-2s: large (>520TB), medium (<520 TB, but significant CPU), and small. The small sites can probably directly access data at another site (but check network connectivity). The medium sites will probably need a disk buffer.
- * Gareth suggested to refer to a "buffer" rather than a "cache". The purpose of the buffer is to prestage files for directio access. It's most efficient if this can be coordinated with job starts, as ARC-CE does (eg. currently at Durham). XCache may still help, but not coordinated with job starts.
- * For small sites with directio over the WAN, need to find another site close by that can handle the load on the storage. This could be a particular problem for Southgrid, which has several small sites.
- * Sam is developing a network map of who is close to whom, though network topologies can change with time.
- * We need better monitoring. Maybe we can get something out of PanDA or Monit, but needs someone to do the research (not a trivial job).
- * Any disk freed up by these moves can be allocated as LOCALGROUPDISK.
- * Sheffield will have to reduce to 150TB, but has 800 job slots. This is a good candidate for going diskless (no buffer to start with), accessing data from Lancaster and/or Manchester.
- * Birmingham should have 600 TB (700-800TB next year). We can use this! It is large enough for a DATADISK.
- * Agreed to discuss with ATLAS at the Sites Jamboree (5-8 March). We can then bring it back to GridPP.