UNIVERSITY OF BIRMINGHAM

Xcache in Birmingham

Xcache has been in production for the last month. The site consists of:

- Basic server (16C, 24GB) to act as redirector
- Two 40TB pool nodes
- ~400 cores for Atlas → 50 simultaneous 8-core jobs
- Pool nodes are connected at 10Gb/s, workers at 1Gb/s.

The XCache instance was setup manually rather than through SLATE because of university policies/concerns over allowing external users direct access to the network

Setup was pretty easy and didn't seem to highlight any issues on the site side that weren't of my own making!

UNIVERSITY OF BIRMINGHAM

Xcache in Birmingham

From monitoring the XCache nodes, I have drawn the following conclusions:

- The caching mechanism seems to work as intended
- New files are cached for arriving jobs then access ~3-5 times soon after
- During production, disks were filled at ~ 2.5TB/day
- From network graphs, it seems there is a some (significant?) saving on external traffic

This system has been pretty easy to setup and run – many thanks for the help from the GridPP and Atlas teams for making this relatively painless!

It's also worth noting the cache is not filled yet so no cleaning has been done!

However, at the very least it is a very easy way to fill a small amount of storage with useful files and therefore allow (almost) diskless sites to run without overloading the larger ones.

Just to mention: In the last 2 weeks I've updated to C7 VAC machines and then hit some aircon problems so the job rate has fallen considerably – this should be sorted in the next few days!

Cache Metrics

Data from single pool node during active production activity (11/6 → 24/6)

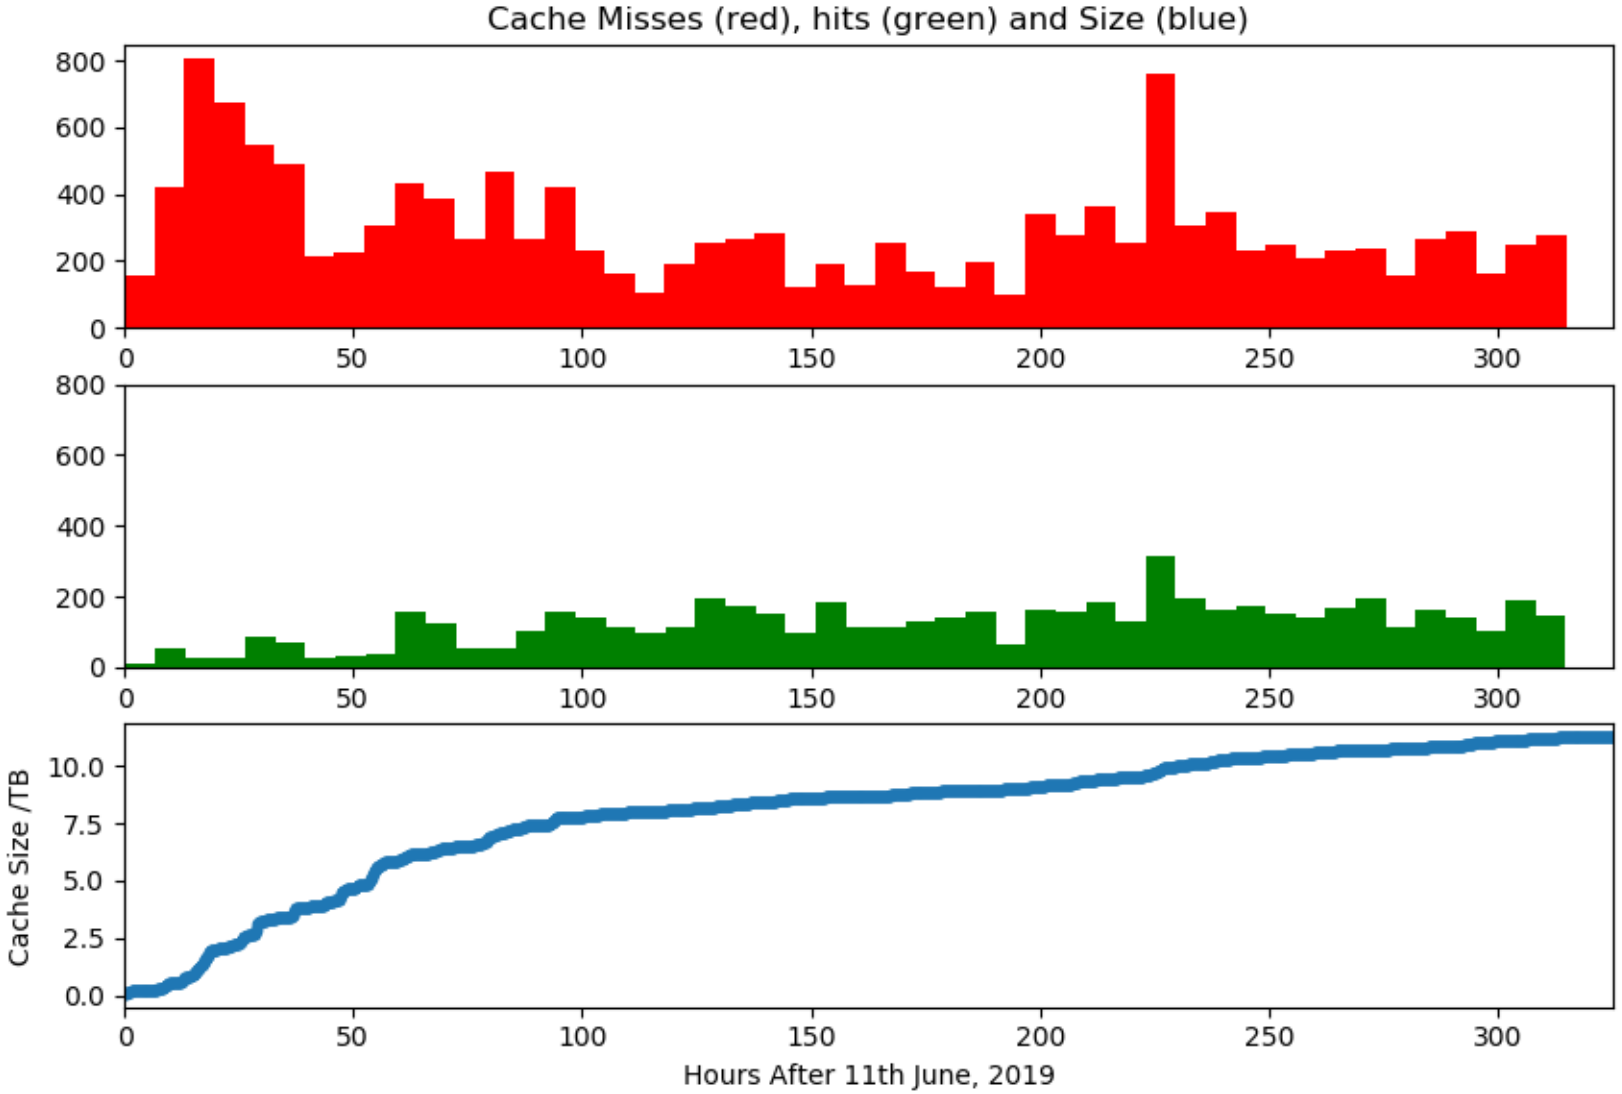

Network Activity

Incoming External and Outgoing Internal Network Activity for a single pool node over a week – Note internal outgoing is significantly larger than external incoming