

Some info I have gathered about the machines in our cluster.

There are 21 compute nodes sitting behind the head node.

All nodes are connected via 1GbE via two switches that share a 6x 1GbE LAG. Exceptions to this are:

--- The head node is connected via x2 1GbE LAG

--- The cluster file server is connected via x6 1GbE LAG

--- For compute nodes compute-0-10 through 0-19, these nodes are equipped with IB running at 40Gb/s (QDR)

All models use identical processors; 2x "Intel(R) Xeon(R) CPU E5-2620 0 @ 2.00GHz", for a total of 12 physical cores and 24 threads per node

Datasheet from Intel: http://ark.intel.com/products/64594/Intel-Xeon-Processor-E5-2620-15M-Cache-2_00-GHz-7_20-GTs-Intel-QPI

The head node has 128G. Compute-0-0 through compute-0-17 have 64G. Compute-0-18 and compute-0-19 have 128G (these are GPU nodes that are moving out of the main pool soon).

The datasheet above didn't give enough info about the cache sizes. This link <http://www.cpu-world.com/CPUs/Xeon/Intel-Xeon%20E5-2620.html> gave the following sizes for the caches. The 6 x is due to the fact that there are 6 cores. So it looks like each processor has a 32K L1 I and D cache and a 256K L2 cache. All cores share a 15M L3 cache.

Level 1 cache size	6 x 32 KB instruction caches 6 x 32 KB data caches
Level 2 cache size	6 x 256 KB
Level 3 cache size	15 MB