

MachineScaling

Among the software questions for setting up and running Hadoop, there are a few other questions that relate to hardware scaling:

1. What are the optimum machine configurations for running a Hadoop cluster?
2. Should I use a smaller number of high end/performance machines or are a larger number of "commodity" machines?
3. How does the Hadoop/Parallel Distributed Processing community define "commodity"?

Note: The initial section of this page will focus on datanodes.

In answer to 1 and 2 above, we can group the possible hardware options into 3 rough categories:

1. database class machine with many (>10) fast SAS drives and >10GB memory, dual or quad x quad core CPU's. With an approximate cost of \$20K.
2. generic production machine with 2 x 250GB SATA drives, 4-12GB RAM, dual x dual core CPU's (=Dell 1950). Cost is about \$2-5K.
3. POS beige box machine with 2 x SATA drives of variable size, 4 GB RAM, single dual core CPU. Cost is about \$1K.

For a \$50K budget, most users would take 25x(B) over 50x(C) due to simpler and smaller admin issues even though cost/performance would be nominally about the same. Most users would avoid 2x(A) like the plague.

For the discussion to 3, "commodity" hardware is best defined as consisting of standardized, easily available components which can be purchased from multiple distributors/retailers. Given this definition there are still ranges of quality that can be purchased for your cluster. As mentioned above, users generally avoid the low-end, cheap solutions. The primary motivating force to avoid low-end solutions is "real" cost; cheap parts mean greater number of failures requiring more maintenance/cost. Many users spend \$2K-\$5K per machine. For a longer discussion of "scaling out" reference: <http://jcole.us/blog/archives/2007/06/10/scaling-out-and-up-a-compromise/>

More specifics:

Multi-core boxes tend to give more computation per dollar, per watt and per unit of operational maintenance. But the highest clockrate processors tend to not be cost-effective, as do the very largest drives. So moderately high-end commodity hardware is the most cost-effective for Hadoop today.

Some users use cast-off machines that were not reliable enough for other applications. These machines originally cost about 2/3 what normal production boxes cost and achieve almost exactly 1/2 as much. Production boxes are typically dual CPU's with dual cores.

RAM:

Many users find that most Hadoop applications are very small in memory consumption. Users tend to have 4-8 GB machines with 2GB probably being too little. Hadoop benefits greatly from ECC memory, which is not low-end, however using ECC memory is RECOMMENDED. see Dennis Kubes' discussion at http://mail-archives.apache.org/mod_mbox/hadoop-core-dev/200705.mbox/%3C465C3065.9050501@dragonflymc.com%3E