

Cassandra Hardware

Memory

The most recently written data resides in memory tables (aka [memtables](#)), but older data that has been flushed to disk can be kept in the OS's file-system cache. In other words, *the more memory, the better*, with 4GB being the minimum we typically recommended in a virtualized environment (e.g., EC2 Large instances). Obviously there is no benefit to having more RAM than your hot data set, but with dedicated hardware there is no reason to use less than 8GB or 16GB, and you often see clusters with 32 GB or more per node.

RAM can also be useful for the key cache (introduced in 0.5) and row cache (in 0.6).

CPU

Many workloads will actually be CPU-bound in Cassandra before being memory-bound. Cassandra is highly concurrent and will make good use of however many cores you can give it. For raw hardware, 8-core boxes are the current price/performance sweet spot. If you're running on virtualized machines, consider using a provider such as Rackspace Cloud Servers that allows CPU bursting.

Disk

The short answer here is that ideally you will have at least 2 disks, one to keep your `CommitLogDirectory` on, the other to use in `DataFileDirectories`. The exact answer though depends a lot on your usage so it's important to understand what is going on here.

Cassandra persists data to disk for two very different purposes. The first is to the commitlog when a new write is made so that it can be replayed after a crash or system shutdown. The second is to the data directory when thresholds are exceeded and memtables are flushed to disk as SSTables.

Commit logs receive every write made to a Cassandra node and have the potential to block client operations, but they are only ever read on node start-up. SSTable (data file) writes on the other hand occur asynchronously, but are read to satisfy client look-ups. SSTables are also periodically merged and rewritten in a process called *compaction*. Another important difference between commitlog and sstables is that commit logs are purged after the corresponding data has been flushed to disk as an SSTable, so `CommitLogDirectory` only holds uncommitted data while the directories in `DataFileDirectories` store all of the data written to a node.

So to summarize, if you use a different device for your `CommitLogDirectory` it needn't be large, but it should be fast enough to receive all of your writes (as appends, i.e., sequential i/o). Then, use one or more devices for `DataFileDirectories` and make sure they are both large enough to house all of your data, and fast enough to both satisfy reads that are not cached in memory and to keep up with flushing and compaction.

As covered in [MemtableSSTable](#), compactions can require up to 100% of your in-use space temporarily in the worst case, free on a single volume (that is, in a data file directory). So if you are going to be approaching 50% or more of your disks' capacity, you should raid0 your data directory volumes. B. Todd Burruss adds on the mailing list, "With the file sizes we're talking about with cassandra and other database products, the [raid] stripe size doesn't seem to matter. Mine is set to 128k, which produced the same results as 16k and 256k." In addition to giving you capacity for compactions, raid0 will help smooth out io hotspots within a single sstable.

On ext2/ext3 the maximum file size is 2TB, even on a 64 bit kernel. On ext4 that goes up to 16TB. Since Cassandra can use almost half your disk space on a single file, if you are raiding large disks together you may want to use XFS instead, particularly if you are using a 32-bit kernel. XFS file size limits are 16TB max on a 32 bit kernel, and basically unlimited on 64 bit.

Cloud

Several heavy users of Cassandra deploy in the cloud, e.g. [CloudKick](#) on Rackspace Cloud Servers and [SimpleGeo](#) on Amazon EC2.

On EC2, the best practice is to use L or XL instances with local storage. I/o performance is proportionately much worse on S and M sizes, and EBS is a bad fit for several reasons (see [Erik Onnen's excellent explanation](#)). Put the Cassandra commitlog on the root volume, and the data directory on the raid0'd ephemeral disks.

<https://c.statcounter.com/9397521/0/fe557aad/1/> | stats