

Cassandra EC2

With these simple bash scripts you can easily start up a small Cassandra cluster on EC2 for testing. [Cassandra EC2](#) See the readme for instructions on how to use it.

_ Note: I found these scripts a useful template but they are a bit dated now and required a variety of fixes, in particular:

- SVN URL has changed from <http://svn.apache.org/repos/asf/incubator/cassandra/trunk/> to <https://svn.apache.org/repos/asf/cassandra/trunk> (no longer in incubation)
- storage-conf.xml is now cassandra.yaml
- The debian (lenny) AML would not apply the updates required due to an issue with the signing key (I don't know the way to fix this, as I'm not a debian person but it's probably trivial to fix)_

Worth noting that this is not intended for production use.

- The data is stored locally on the EC2 instance and will disappear when it is terminated.
- EC2 io performance is often not great compared to non-virtualized architectures. You should balance the rest of your hardware allocation accordingly.

The [Cluster Chef](#) repo has chef scripts and poolparty recipes that can launch and provision a scalable cassandra cluster in the EC2 cloud. It can create either a persistent cluster (using EBS volumes that remain when the instances are decommissioned) or an ephemeral cluster (that uses the local scratch drives), and has optional support for NFS and hadoop. The chef cookbook is based on one by Benjamin Black

Edit 11-9-11

- The [DataStax AML](#) provides a quick and easy way to get Cassandra 1.0 up and running in minutes.
- The [DataStax Chef](#) repository provides yet another set of scripts for installing Cassandra 07x and 08x. These are updated a bit behind release schedule.

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