

## Laboratory #2

Released: October 11, 2016

Due: October 21, 2016 (Friday)

In this lab, you will explore volatile virtual machines computing as a resource models, and measure some basic dimensions of performance, including launch time, cost, and performance. You have the choice picking two infrastructures that provide this service (Amazon's spot instances and Google's preemptible VM's), and I hope you will as a class spread to cover the options. In addition to turning in a writeup, that plots and graphs your results appropriately, you will give a short presentation to the class on what you learned about the platform from the experience.

1. (get started) Pick one of the two preemptible virtual machine cloud implementations
  - a. Amazon
    - i. Watch the tutorials and read documentation. Amazon: <https://aws.amazon.com/ec2/spot/spot-tutorials/>
    - ii. Experiment with creating spot requests and terminating spot instances using the console
    - iii. Be careful to clean up properly! Cancelling a request does **NOT** terminate the spot instance (this is a bit odd, and always causes problems)
  - b. Google
    - i. Watch the tutorials and read documentation. Google: <https://cloud.google.com/preemptible-vms/>
    - ii. Experiment with creating preemptible VM's and terminating them using the console
2. (experiment with volatility) In this part, your objective is to experiment with volatility in spot instance and preemptible VM's. This is a little easier to do with Spot instances, where you can manipulate this with low bids, but also possible with Google VM's. (Hint: you probably want to use the command-line interface (CLI) or one of the other interfaces (Python, Java, etc.) and write some programs to do this systematically. And, the security configuration can be a little tricky to get right, so read some of that documentation and test carefully before you scaleup.
  - a. Amazon (before you start, spend some time looking at these pricing sites: <http://predictspotprice.cs.ucsb.edu> and <http://ec2price.com> and use this information both to pick what instance types you will do experiments on, as well as the bidding levels.
    - i. Design an experiment that uses one of the cheapest spot instances (c1.small, t1.micro, perhaps), and explores its volatility across times of day and regions. Starting from well below average of market price (this may be below the current market price), and working up to a fraction of on-demand price (75%), design a set of experiments that explore the usability of spot instances. It might be helpful to look here <https://aws.amazon.com/ec2/spot/bid-advisor/>, but be careful about using any of the more expensive instances.
    - ii. Vary the bid price for spot instances across this range, how does the time to launch an instance depend on bid price? Once they've launched, let them run for awhile (let's cap these at 24 hours). How long do they run without interruption, again note how this varies as a function of the bid price?
    - iii. Look for awhile at several different regions in <http://ec2price.com/> and find a region or AZ that for your cheap instance type is significantly different

- from the one you studied. Repeat the same set of experiments again, and document the results.
- iv. Compare the results from the two regions/AZs, how are they similar, how are they different? How would these differences affect the cost or usability of spot instances in those regions/AZs? In late 2015, Amazon introduced defined duration spot instances. How does this compare to what you measured with spot instances?
- b. Google
    - i. Here the situation is a bit simpler, as the Google cloud has no bidding system. Pick one of the cheaper preemptible VM types, and perform a set of experiments that launch preemptible VM's
    - ii. Launch preemptible VM's, at varying times of day over a 24 hour period (spread over a longer period if you can manage it). Once they've launched, let them run for awhile (let's cap each at 24 hours). How long do they run without interruption? What does the distribution look like? Can you find any structure?
3. (computation with a cluster)
- a. Amazon: Starcluster or Hadoop
    - i. Pick one of these two systems (or an alternative cluster use system), and use spot instances to increase performance on some simple computations (demonstration applications for these platforms, for example). Be sure to use pretty small runs!
    - ii. Setup some infrastructure nodes using on-demand instances, and use the spot instances to scale computing up – perhaps up to 4 additional small VM's. The goal here is to do a small experiment, but show a nice linear speedup.
  - b. Google
    - i. Find a working cluster management system for spinning up clusters in Google Compute Engine. Make sure that it can drive preemptible instances. Use preemptible instances to increase performance on some simple computations (demonstration applications for these platforms, for example). Be sure to use pretty small runs!
    - ii. Setup some infrastructure nodes using on-demand instances, and use the preemptible instances to scale computing up – perhaps up to 4 additional small VM's. The goal here is to do a small experiment, but show a nice linear speedup.