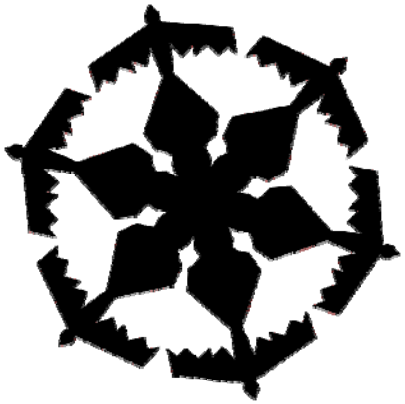

Five Minutes of Rage with SnowFlock

H. Andrés Lagar-Cavilla
andreslc@cs.toronto.edu
Joe Whitney, Adin Scannell, Steve Rumble, Philip Patchin,
Mike Brudno, Eyal de Lara, M. Satyanarayanan
University of Toronto (mostly)

http://sysweb.cs.toronto.edu/snowflock

The Basics

- Cloud Computing
- Virtualization

If you are not a believer, then you can

The Basics

- Cloud Computing
- Virtualization

If you are not a believer, then you can

- Resume your email checking activities
- Return to your placid sleep
- Dedicate to the noble art of schmoozing
- Desert the room in search for caffeine

The Rant

- When I signed up for the cloud
- The ad said: “100s of CPUs at your fingertips”
- What it didn’t said was that:
  - New VMs show up in minutes
  - State needs to be pushed explicitly

New VMs Show Up in Minutes

- #@☠☹💣! I missed my flash crowd…
- I wait two mins to start my 40sec parallel task
  - Sequence alignment, rendering, quant pricing
  - One hour w/one processor
  - 40 seconds w/128 processors
- Work around: pre-allocate
  - $$$
  - Consolidation harms performance

State Has to Be Pushed Explicitly

- As in, tell the new guys where we're at
  - Because they're basically fresh boots
  - Or were pre-allocated two weeks ago
- Message passing, etc

Wouldn’t it be nice if the cloud was UNIX?

- You could fork your VMs
- Like processes do
- And wouldn’t it be nice to fork instantaneously
- And make your VMs stateful too…
- Wouldn’t it be nice?

The Beach Boys Were Dreaming of SnowFlock

- Which, through a melange of cunning hacks
- Clones a VM into dozens of identical copies
- In less than a second
- Each executing on a different box
- With negligible runtime overhead
- Thus enabling hundred of CPUs instantly at your fingertips
- As advertised

And As Easy To Program As

```
while (forever)
    receive new request
    fork worker to handle it
    join worker when done

while (forever)
    if webserver load is high
        fork some new workers
    else if webserver load is low
        kill some of those extra workers
        (because I can come up with more workers later in no time)
```

Free And Open Source

- http://sysweb.cs.toronto.edu/snowflock
- H. Andrés Lagar-Cavilla
- andreslc@cs.toronto.edu
- Big team @ University of Toronto