Availability in the Cloud

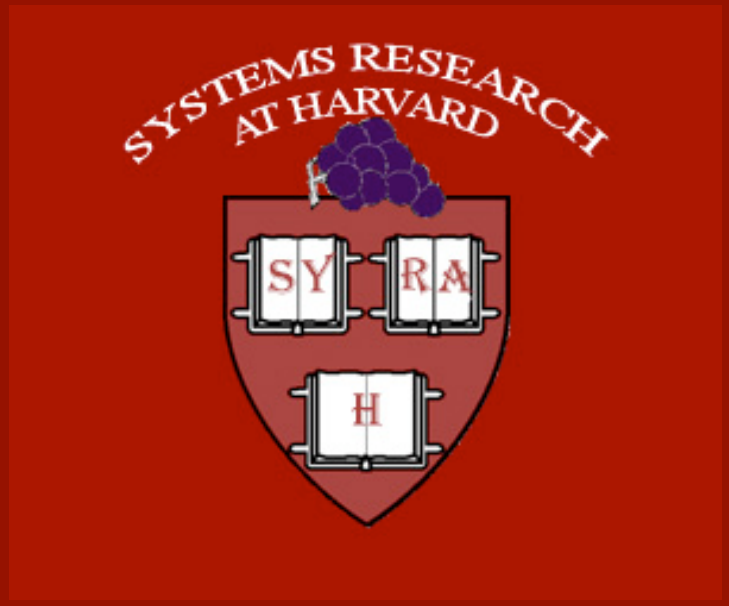

Kiran-Kumar Muniswamy-Reddy

Overview

- Cloud has made acquiring infrastructure easier
- Still need to be a 4-star wizard to build distributed apps
- Goal: Allow a 1-star wizard to build distributed apps

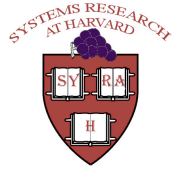

Existing Services (1)

- Google AppEngine
  - Automatic load balancing
  - Persistent state on Megastore
- Works for a narrow set of applications

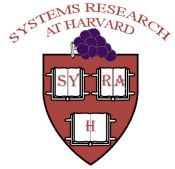

Existing Services (2)

AWS and Azure

- Provide compute and a variety of storage services
  - Also provide a load balancer, cloudwatch

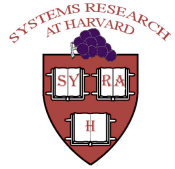

Research Solutions (Transactional Memory)

- Sinfonia – A framework for building infrastructure applications [SOSP'07]
  - Provides a shared address space for nodes to co-ordinate access
  - Data is referenced by specifying (memory-node, address)

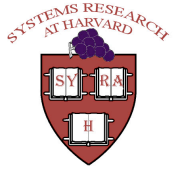

Research Solutions (Transactional Memory)

- Cloud-TM – A distributed transactional memory solution geared for the cloud [LADIS'09]
  - Work in progress
  - Updated their university software to use this

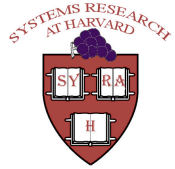

Research Solutions (Language Based)

- Fluxo – provides a dataflow programming model to represent a service's behavior [HotOS'09]
  - Compiles out a deployable program
  - Takes care of partitioning and tiering, fault tolerance etc.
- BOOM – High level language for building distributed systems
  - Declarative language based on Prolog

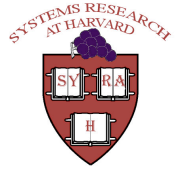

Commercial Solutions

- ISIS toolkit
  - Used by the stock exchanges, air traffic control systems, etc.
  - Process group communication structures
  - Later generation: Horus/Ensemble

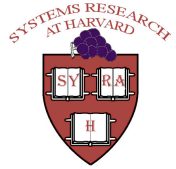

Cloud

- Provides fewer guarantees
  - Inter-node latency
  - Consistency
  - Availability
- Availability across vendors
- Cost is a factor

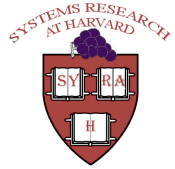

Conclusions

- Cloud has made acquiring infrastructure easier
- To harness its true potential, we need make building apps easier
- What is the right approach?
  - Toolkit, Languages, Transactional Memory or <your favorite approach here>

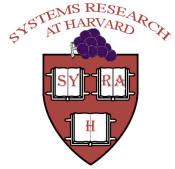

Questions?

Contact:

kiran@eecs.harvard.edu

www.eecs.harvard.edu/~kiran

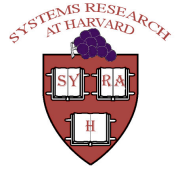

Summary

- Distributed Computing: “The failure of a computer you didn't even know existed can render your own computer unusable” – Leslie Lamport
- Cloud Computing: “Someone you didn't even know existed can render your computer unusable” – Jim Waldo

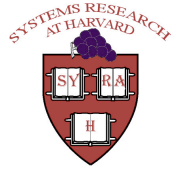

Example App

- Build a better C-store: Harvard C-Store
  - Differentiator: better data layout
- Lots of modules in common
  - Request Processing
  - Query Engine