

Experiences With RAMCloud Applications


John Ousterhout, Jonathan Ellithorpe, Bob Brown



Overview

- **January 2014: RAMCloud 1.0 (first practical version)**
- **Application experiments so far:**
 - Stanford: natural language processing, graph algorithms
 - Open Networking Laboratory: ONOS (operating system for software defined networks)
 - CERN: high energy physics (visiting scientist, summer 2014)
 - Huawei: real-time device management
- **Challenges**
 - Low-latency networking not yet commonplace
 - RAMCloud not cost-effective at small scale
 - RAMCloud is too slow (!!)

Scale and Recovery




Cluster Size	Server Capacity	Cluster Capacity	Recovery Time
101 servers	50 GB	5 TB	1 sec
201 servers	100 GB	20 TB	1 sec
6 servers	100 GB	600 GB	40 sec
6 servers	2.5 GB	15 GB	1 sec
11 servers	5 GB	55 GB	1 sec


Small clusters can't have both fast recovery and large capacity/server

Fast But Not Fastest

Choice #1:
5-10 μ s remote access



Choice #2:
50-100ns local access



- **Choice #2 is 100x faster than RAMCloud**
 - And, can store data in application-specific fashion
 - But, data must partition
 - What about persistence?

Conclusion

- **Technology transfer is a numbers game:**
 - Odds of success with any one group or project are low
 - Must try many experiments to find the right fit
- **Our goals:**
 - Learn something from every test case
 - Keep improving RAMCloud
- **Application issues suggest new research opportunities**