

# **RAMCloud: Scalable High-Performance Storage Entirely in DRAM**

New research project at *Stanford*

John Ousterhout, Christos Kozyrakis, David Mazières, Aravind Narayanan, Diego Ongaro, Mendel Rosenblum, Stephen Rumble, and Ryan Stutsman

# Introduction

- This project asks the question "*Are we nearing the end of the line for disk-based storage systems and, if so, **what is next?***"
- The future of datacenter storage?
  - *A **RAMCloud** is a new class of storage intended for datacenters, where all information **lives at all times in fast DRAM.***
- Create **large-scale storage systems** entirely in DRAM
- Interesting combination: **Scale, Low Latency**
- Enable new applications? RAMClouds may enable a new breed of data-intensive applications.

# Motivation

- **The problem:**
  - **Large-scale Web applications** are finding it **increasingly difficult** to scale disk-based systems to get the **performance** they need.
  - Disk technology evolution has favored capacity **over access speed**.
- **The solution:**
  - Shift the home of online data from disk to **DRAM**.
  - Large-scale systems are created by **aggregating** the **main memories** of **thousands of commodity servers**.
  - Combine **large scale** (**100-1000 TBytes**) with **100-1000x faster latency** than current systems (**5-10  $\mu$ s** to access small amounts of RAMCloud data from application servers in the same datacenter).
  - **Simplifying** the creation of **large-scale Web applications**; RAMClouds may enable a new breed of data-intensive applications.
  - **Creating a production-quality implementation of RAMCloud.**
  - *Google and Yahoo! store their search indices entirely in DRAM.*
- Numerous research & challenges:
  - **Durability** of RAMCloud data; **Very low latency**; **Size and Complexity**

# Contributions

- The **RAMCloud concept** and the **scale** and **performance** that are believed is achievable in a RAMCloud system.
- Two motivations for RAMClouds: **Applications** and the underlying **storage technologies**
- Benefits of **extremely low latency** and compares RAMClouds with two alter-native approaches (caching and flash memory).
- Numerous **research challenges** to build a practical RAMCloud will need to be addressed.
- **Disadvantages** of RAMClouds, such as high cost/bit and high energy usage.

**Question?**