# RAMCloud

## Scalable, Low-Latency Datacenter Storage

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

## 100-1000TB at 5-10µs latency

### The Basic Idea

- Storage for datacenters
- 1000-10000 commodity servers
- 32-64 GB DRAM/server
- All data always in RAM
- Durable and available
- Performance goals:
  - High throughput: 1M ops/sec/server
  - Low-latency access: 5-10µs RPC
- Data model similar to key-value store

### Why Latency Matters

Traditional Application: Single Machine

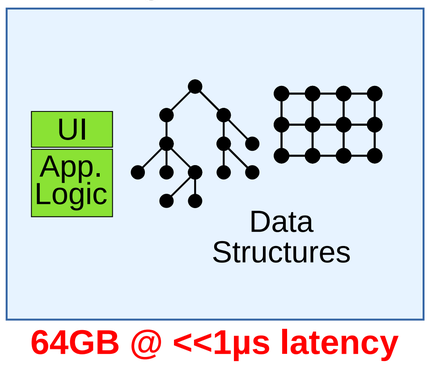


64GB @ <<1µs latency

Web Application: Datacenter

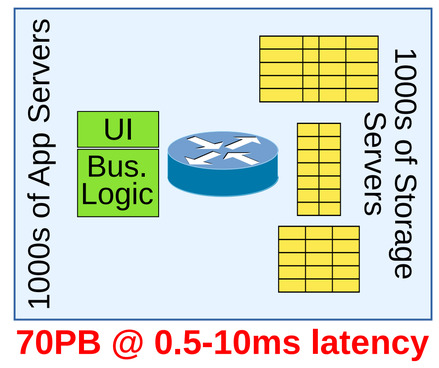


70PB @ 0.5-10ms latency

RAMCloud's goal: best of both worlds, low latency *and* large scale

### Example Configurations

| | Today | 5-10 years |
|---|---|---|
| # servers | 2000 | 4000 |
| GB/server | 24GB | 256GB |
| Total capacity | 48TB | 1PB |
| Total server cost | $3.1M | $6M |
| $/GB | $65 | $6 |

Compare with Facebook:
4000 MySQL servers and 2000 memcached servers
200TB of non-image data

### Cluster Organization

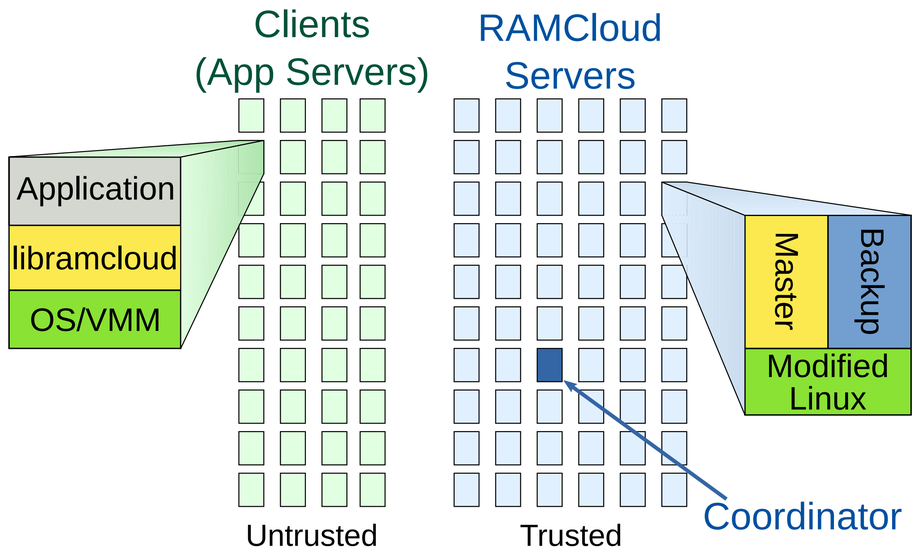


### Durability

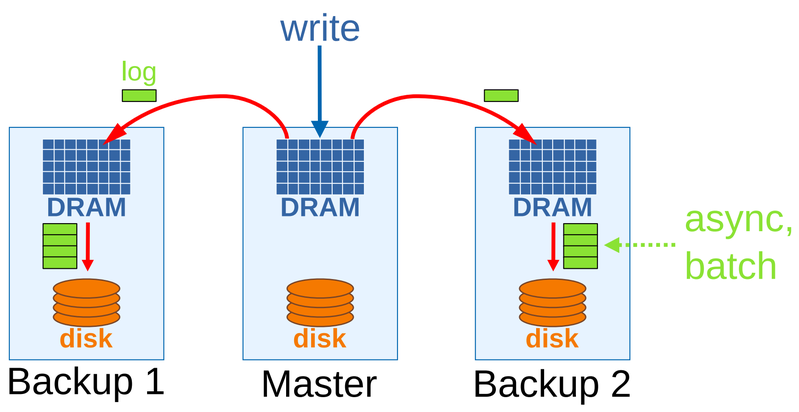


Data is synchronously replicated to *R* backups, then asynchronously written in batch to disk.

### Latency Today

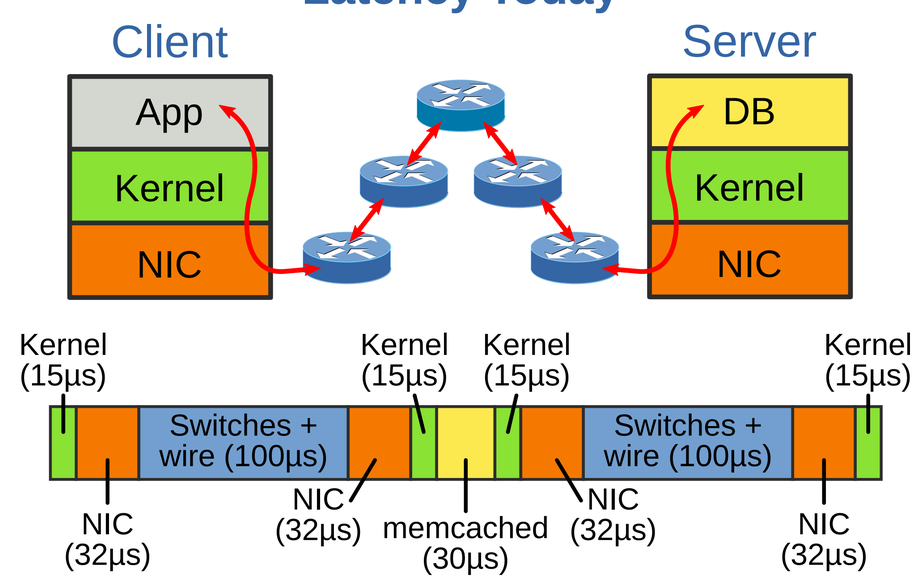


Our tests show 11µs RTT possible (with no switches)