

RAMCloud: Scalable High-Performance Storage Entirely in DRAM

Kevin Chang

15-799 Lightning Talk

10/14/2013

J. Ousterhout et al., "The Case for RAMClouds: Scalable High-Performance Storage Entirely in DRAM", SIGOPS, 2009

DRAM in Storage System

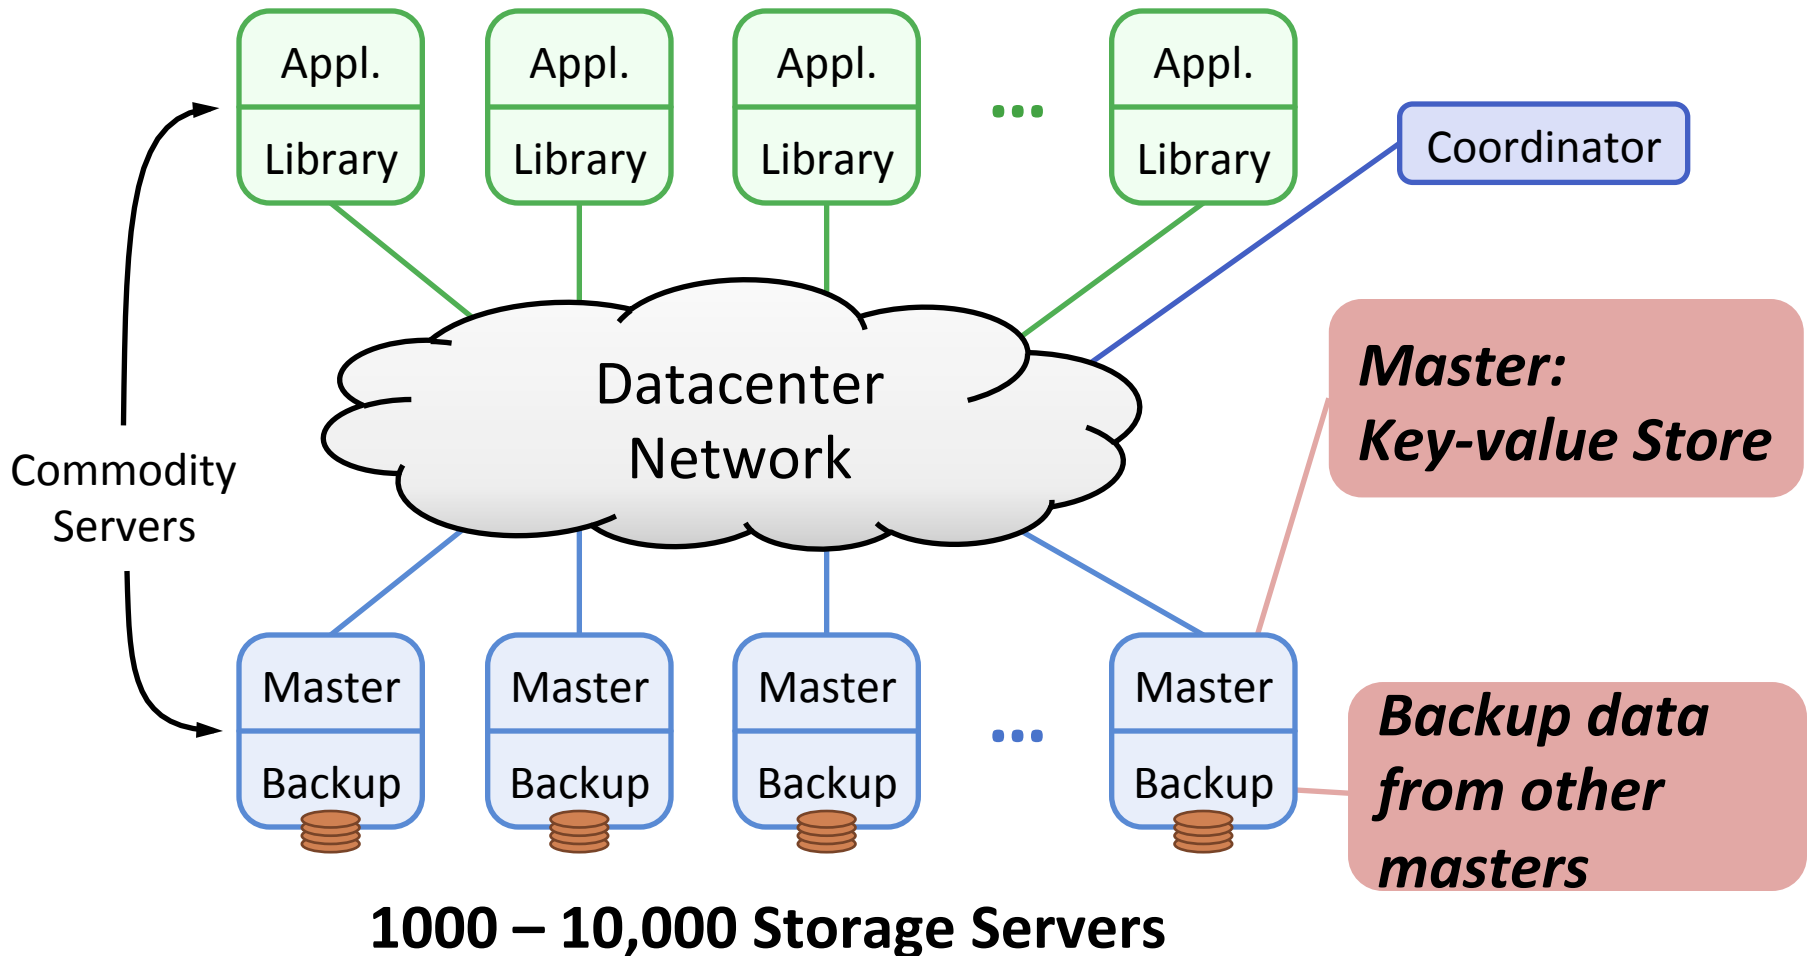
- DRAM has becoming the main storage recently
- Examples:
 - Facebook: 200TB total data, 150TB DRAM cache in 2009
 - Large Web indexes used by Google and Yahoo!
- **Problems**: DRAMs are not efficiently used
 - Typically used as a cache: large penalty from \$misses
 - Synchronous I/O writes with the backing store

RAMCloud

- General-purpose DRAM-storage system
- **Goal**: maximize DRAM performance
- **Features**:
 - All data always in DRAM (no \$ misses)
 - Durable and available
 - Large scale: 1000+nodes, 100+TB
 - Low latency: **5-10 μ s** access to data
 - Enable new class of applications:
 - Large-scale graph algorithms
 - Facebook limits off-box requests to b/w 100-150 to ensure low latency


RAMCloud Architecture

1000 – 100,000 Application Servers



Durability and Availability

- **Goal**: high performance (not limited by disk)
- **Problems**:
 - 1. Synchronous disk writes (slow)?
 - Log-structured memory to buffer writes
 - 1 copy of data in DRAM & multiple copies in disk/flash
 - 2. What if the master server fails?
 - Fast crash recovery @ 64GB/sec
 - Divide data into partitions and distribute across recovery masters



Conclusion

- High-performance storage system
 - **5-10 μ s** latency and 100+TB capacity
- Durability and availability
 - Fast crash recovery and log structures