


The future of datacenter storage is NOT disk, NOT flash, but DRAM: RAMCloud

N. Jayakumar

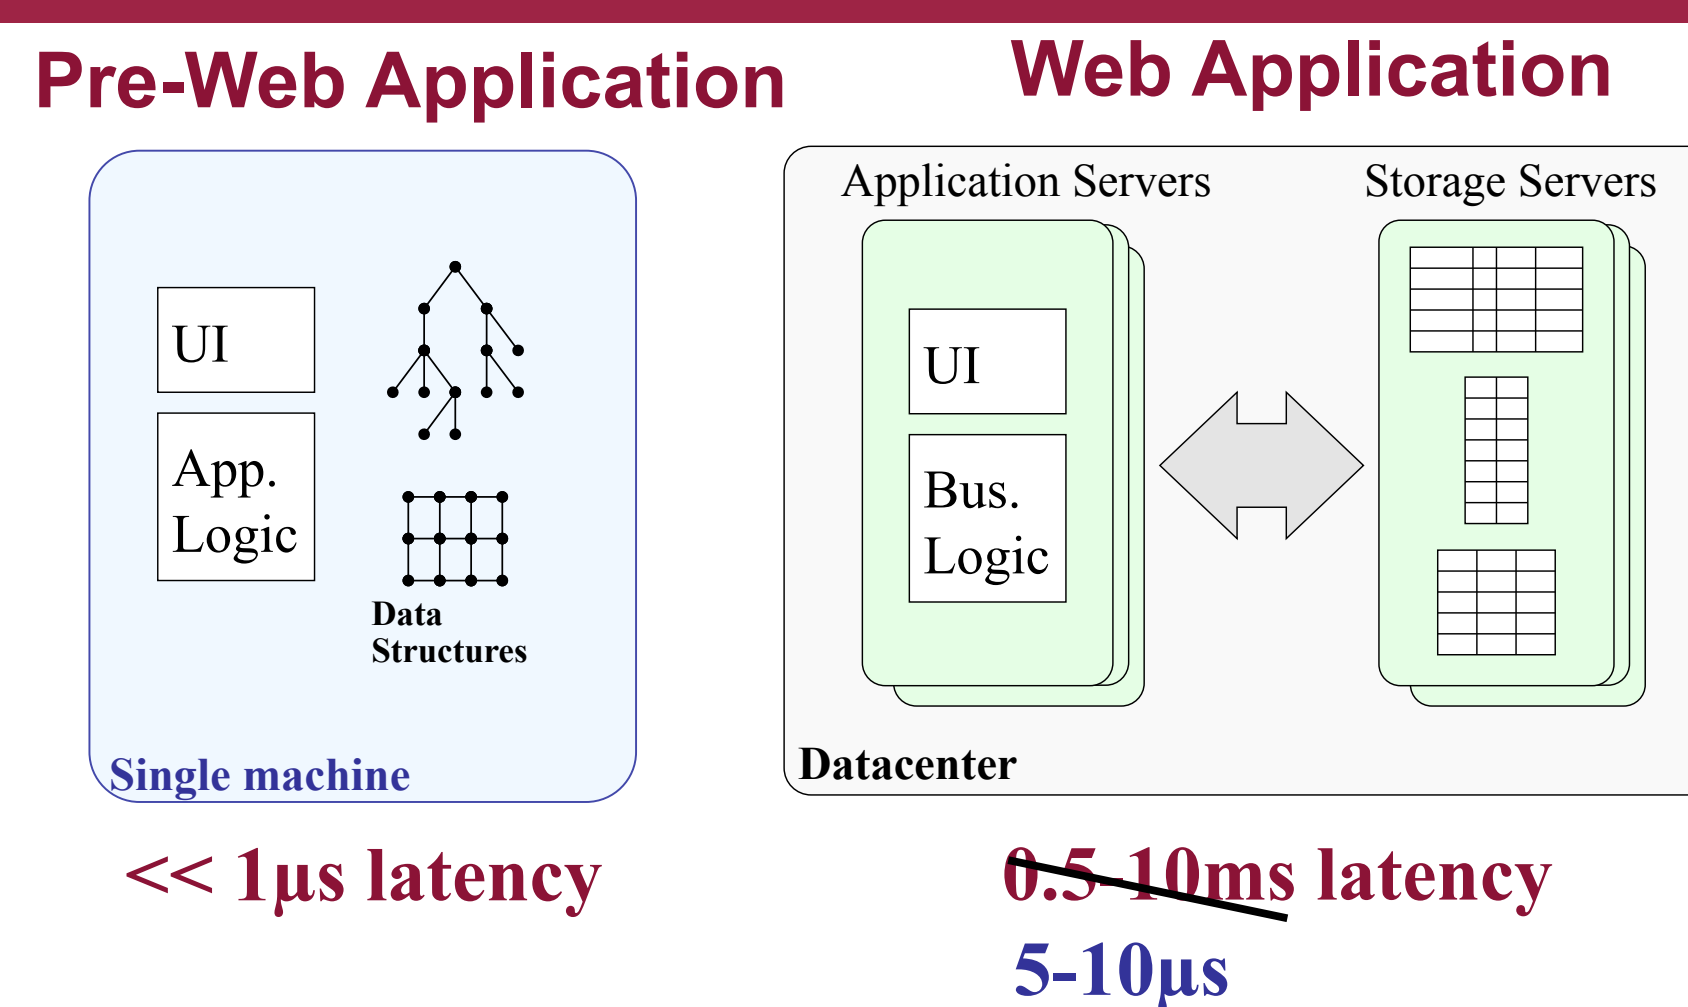
with D. Ongaro, S. Rumble, R. Stutsman, J. Ousterhout, M. Rosenblum
Stanford University



Approach

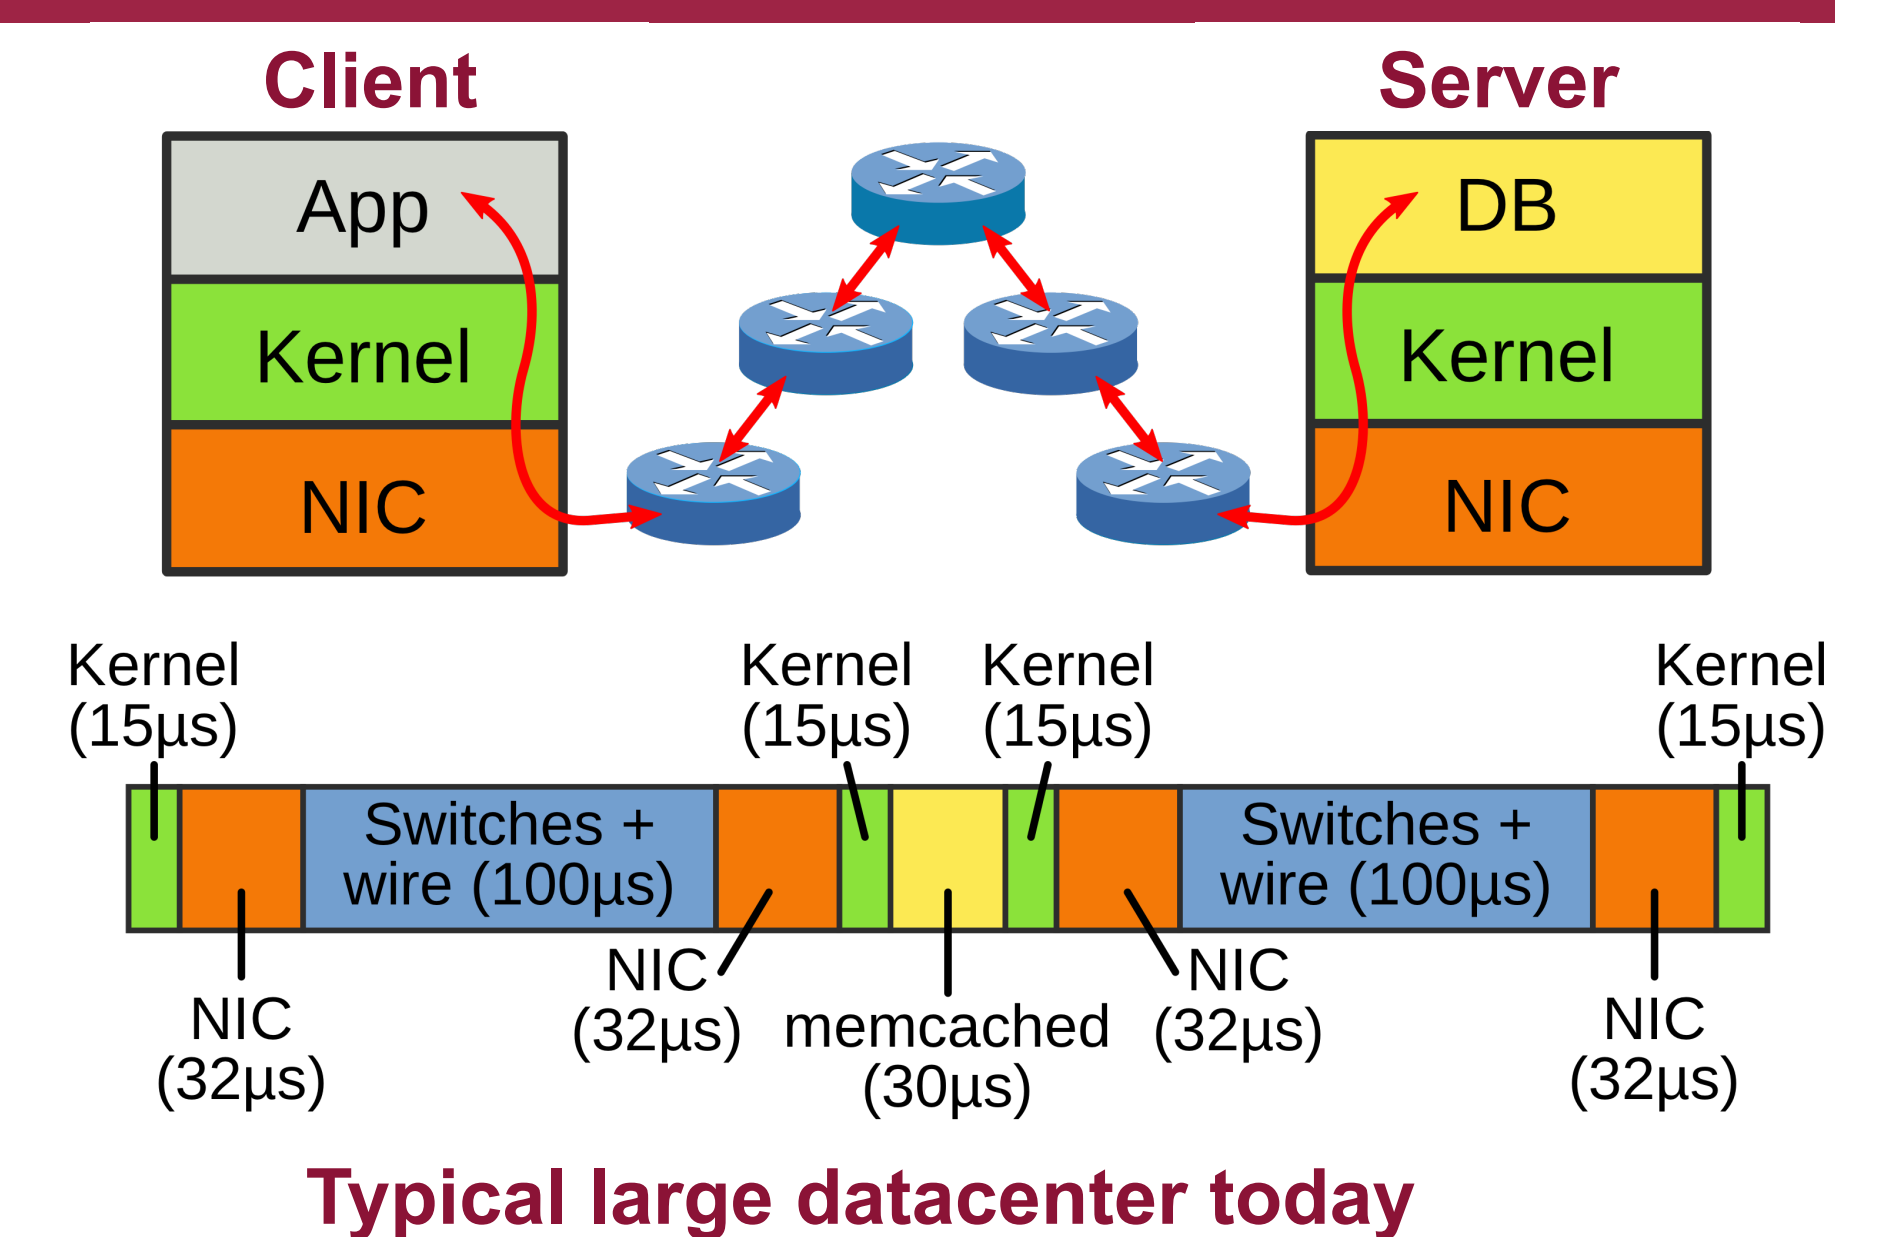
- Storage for datacenters
- 1,000-10,000 commodity servers
- 32-64 GB DRAM/server
- All data always in RAM
- Durable and available
- Performance goals
 - **Throughput: 1M ops/sec/server**
 - **Latency: 5-10 μ s RTT**
- Data access model – Key-Value Store

Why Latency Matters



- RAMCloud goal: large scale **and** low latency
- Enable a new breed of information-intensive applications

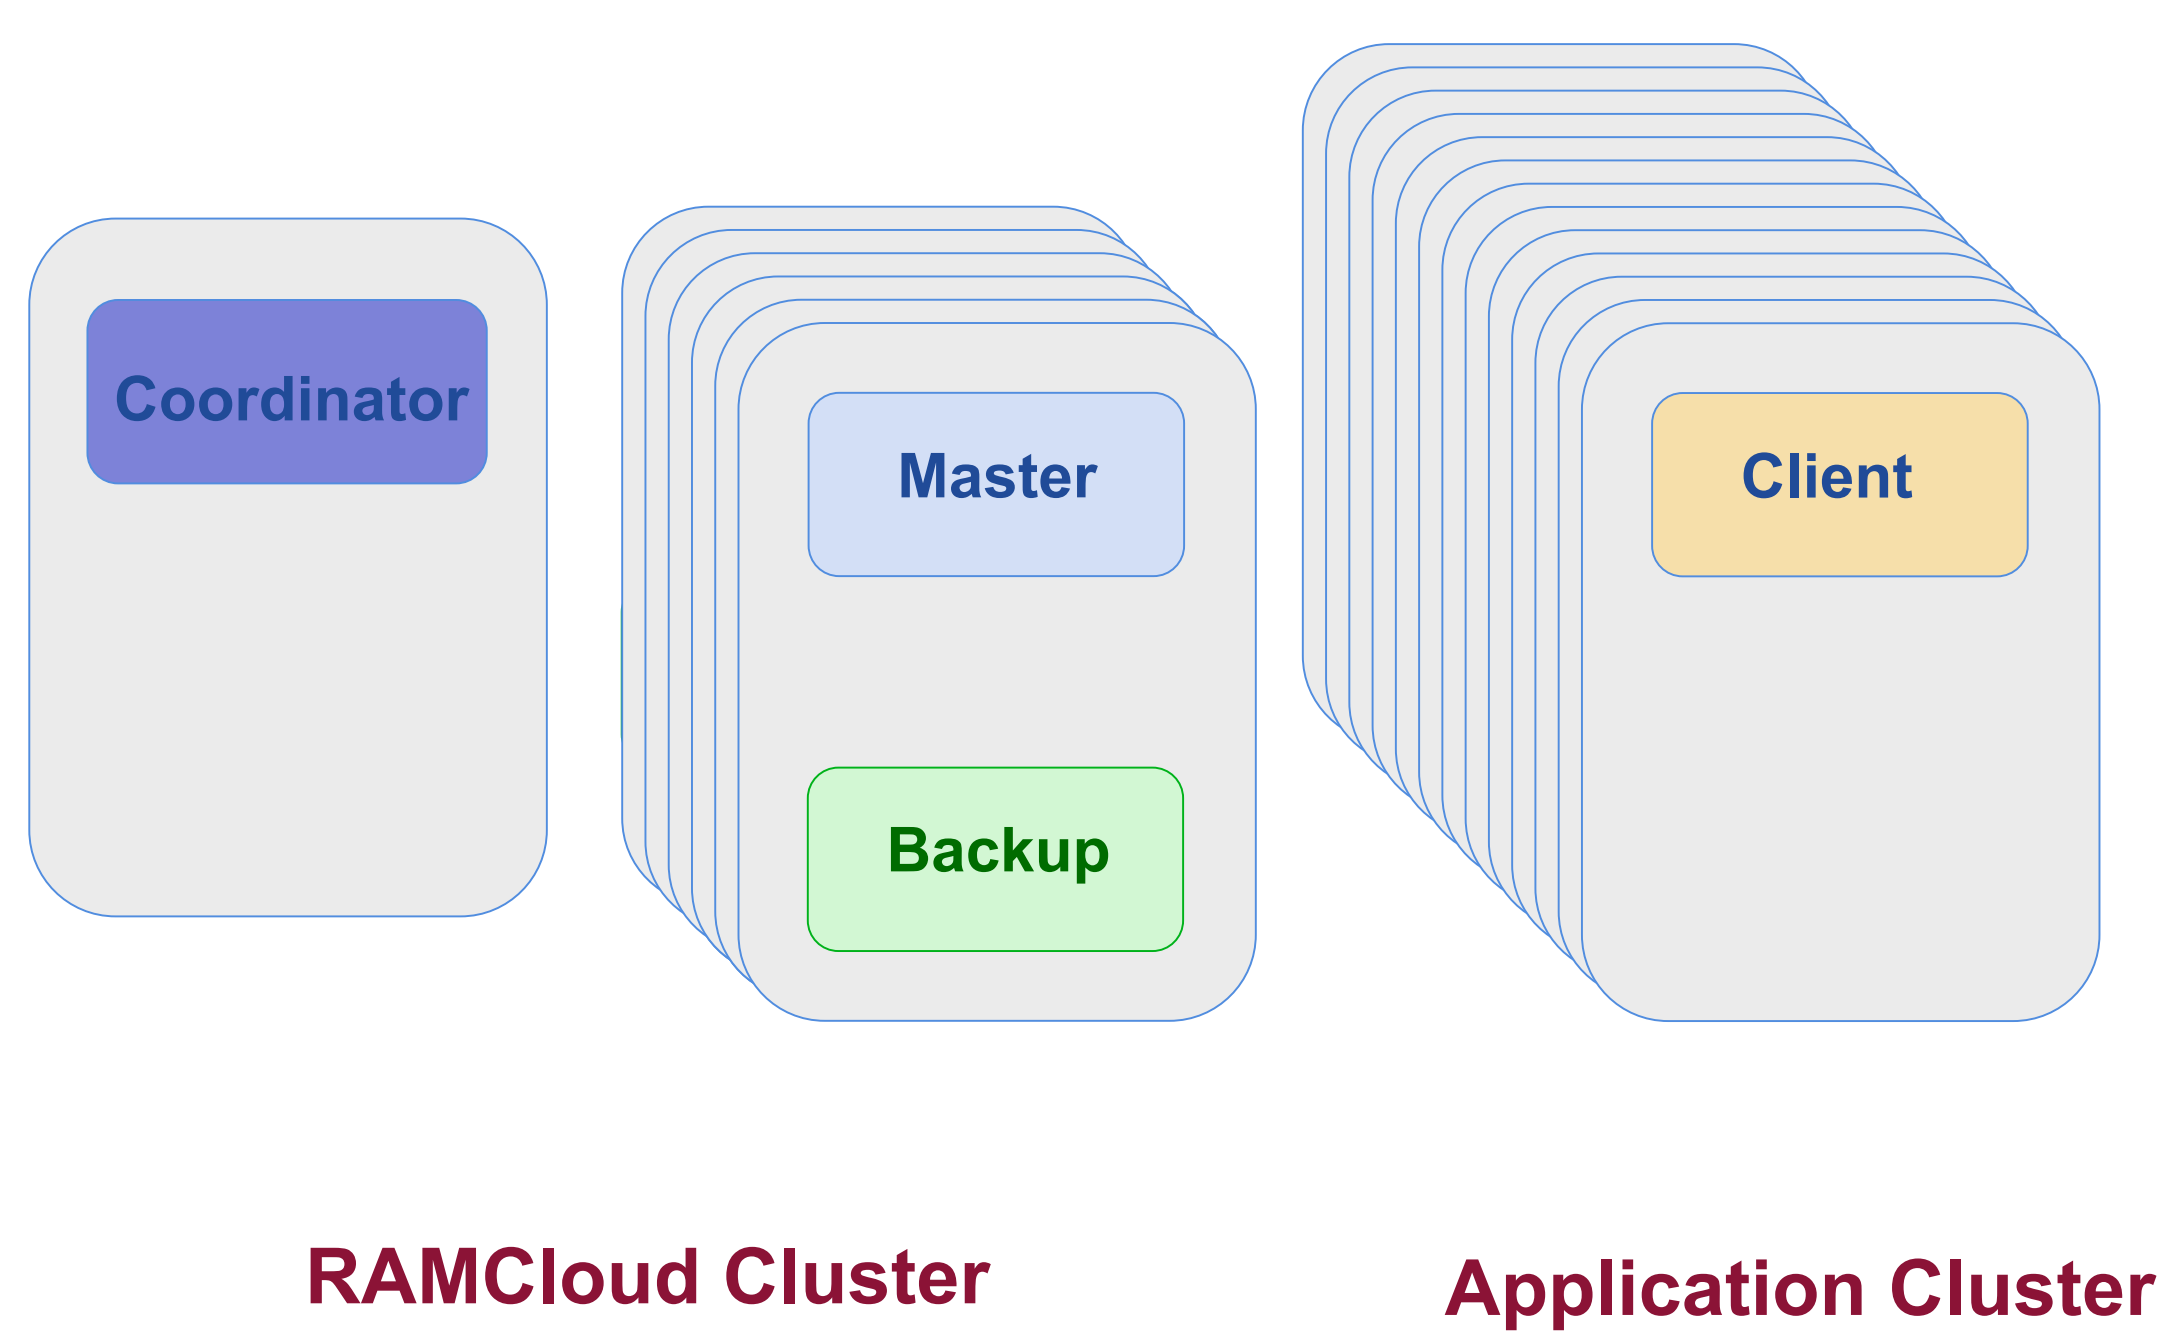
Challenge: RPC Latency



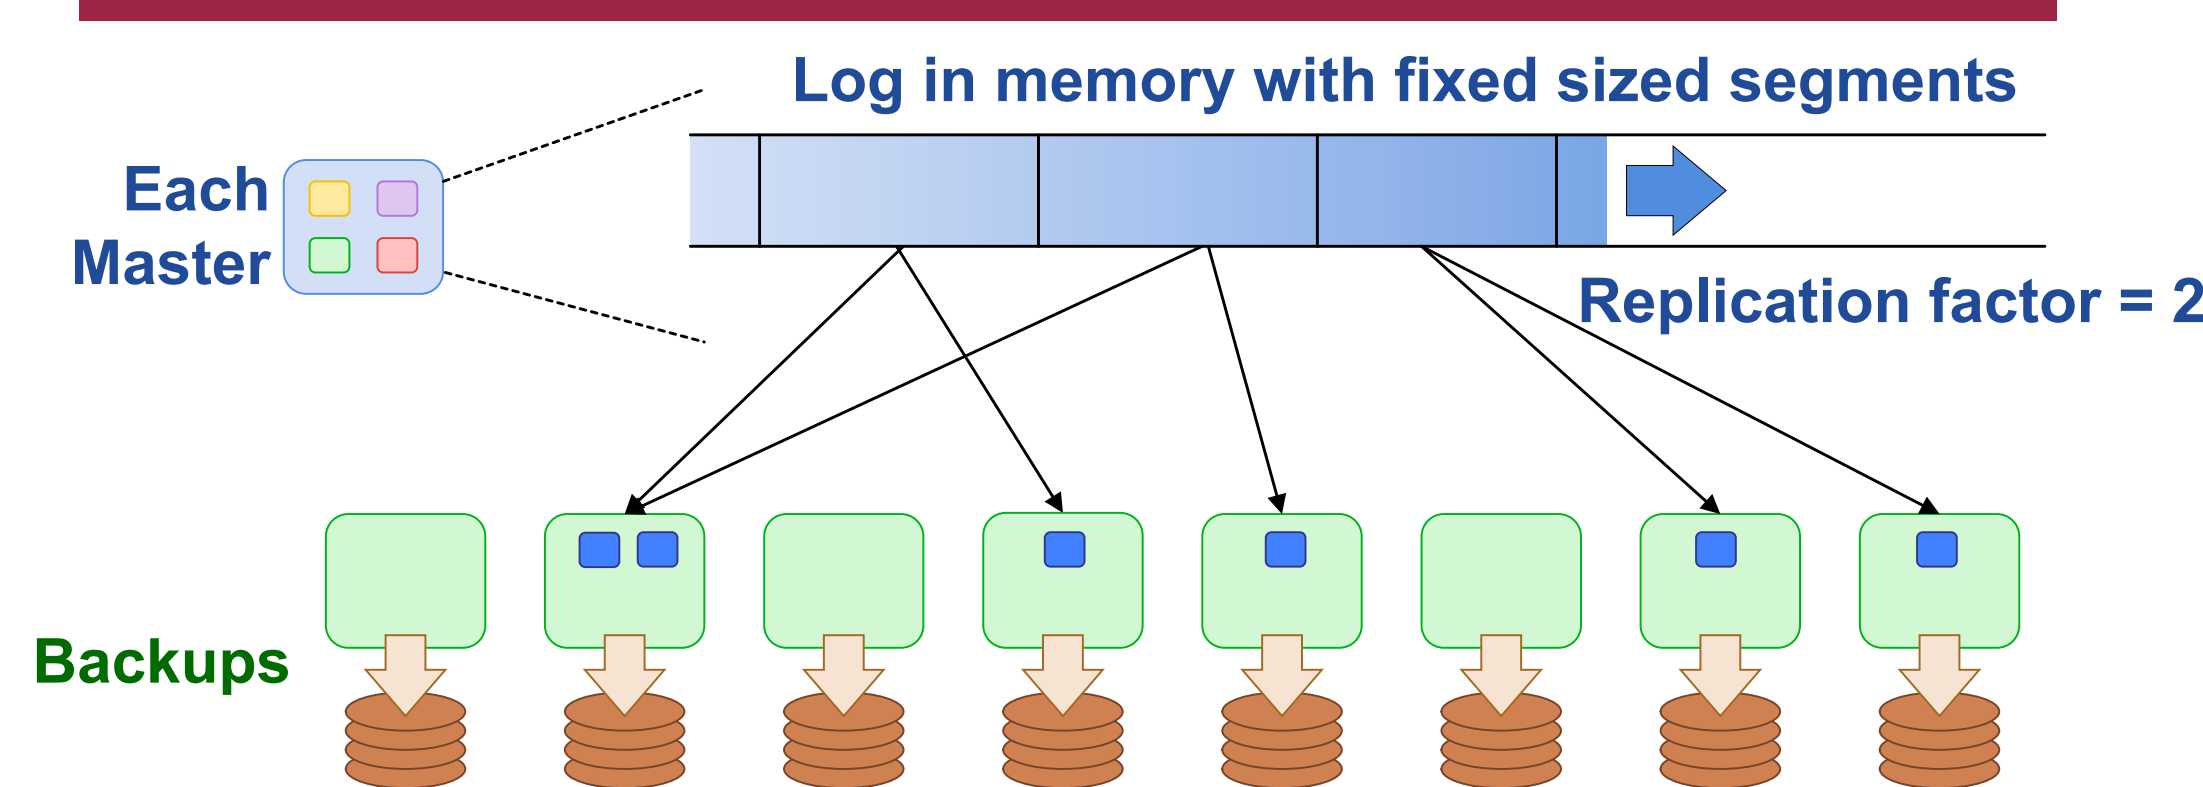
Typical large datacenter today

Our best time: $\sim 7\mu\text{s}$ RTT (Infiniband, 1 switch)

Datacenter with RAMCloud

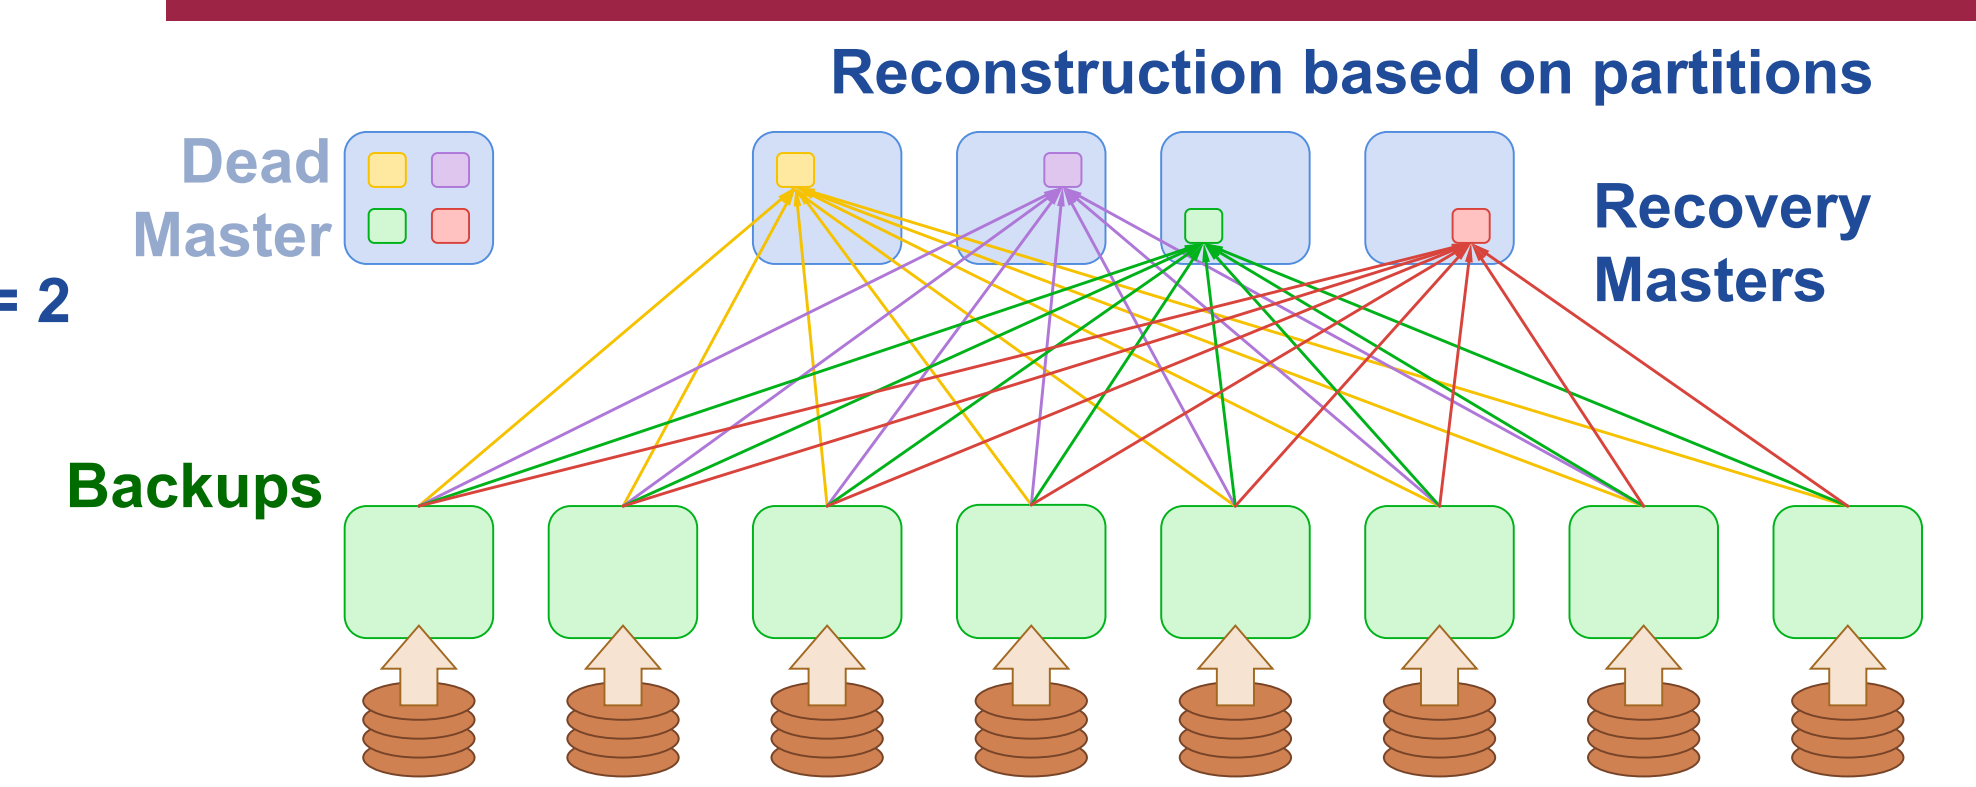


Log-structured Memory



- Data is maintained in a log structure in RAM
- One log per master divided into segments
- Results in optimized disk writes
- Segments are sent to disk (backups) on other nodes chosen randomly.
- Single copy of data in RAM. Redundant copies on disk.

Challenge: Availability



- Fast recovery on node failure as the model for high availability.
- Goal recovery time **1-2 seconds** for node failures
- Divide dead master's data into partitions: each reconstructed on separate masters
- Partitions based on tables & key ranges, not log segments