


Load Balancing in Ceph: Load Balancing With Pseudorandom Placement

Esteban Molina-Estolano, Carlos Maltzahn, Scott Brandt
University of California, Santa Cruz

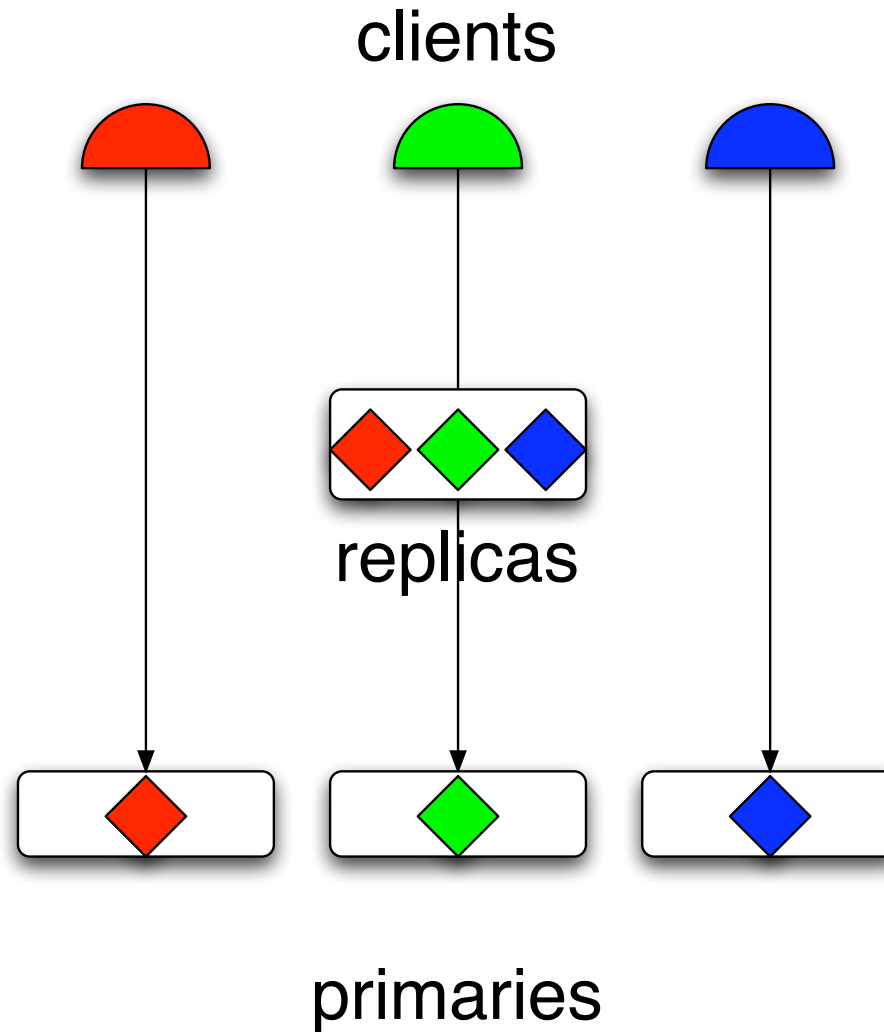
The Load Balancing Problem

- Pseudorandom placement for distributed storage systems has several advantages
- Problem: pseudorandom placement makes load balancing harder
- Research platform: Ceph, object-based storage system developed at UCSC

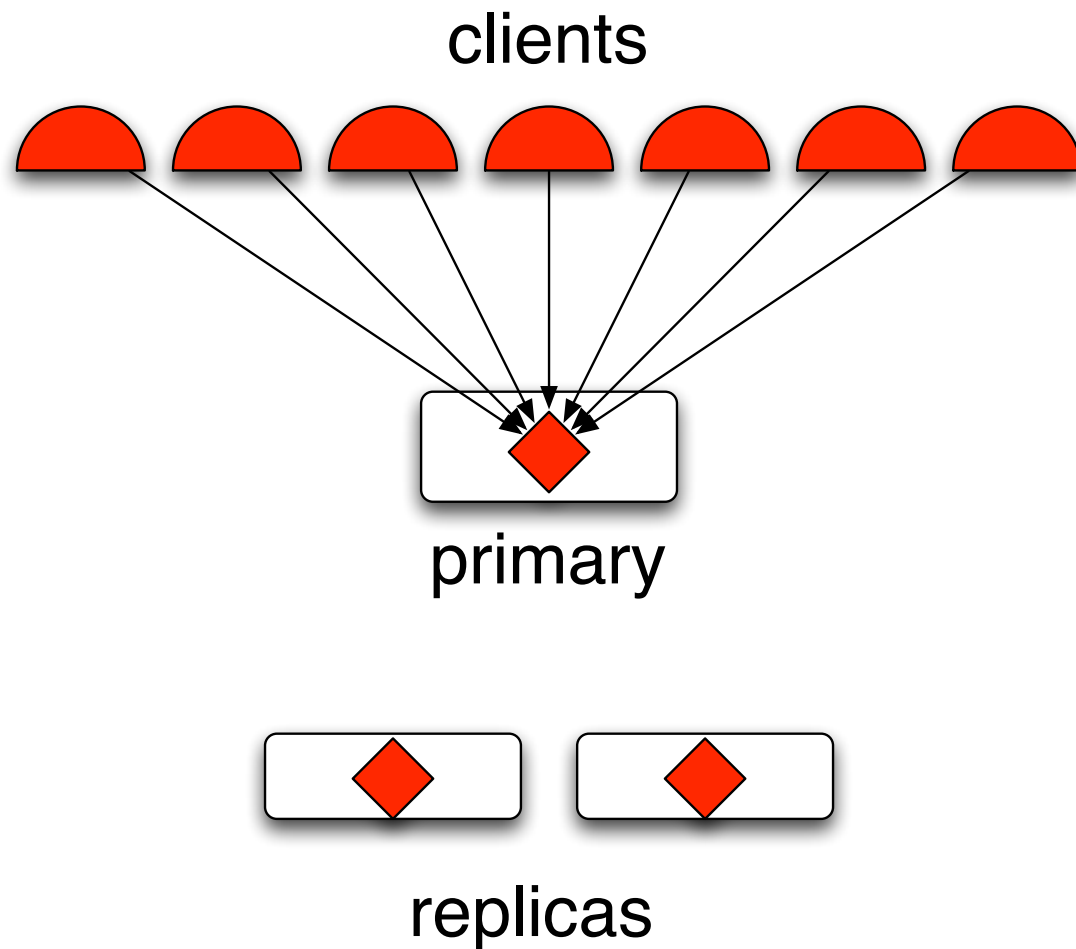
Coincidental overload




Primary switching



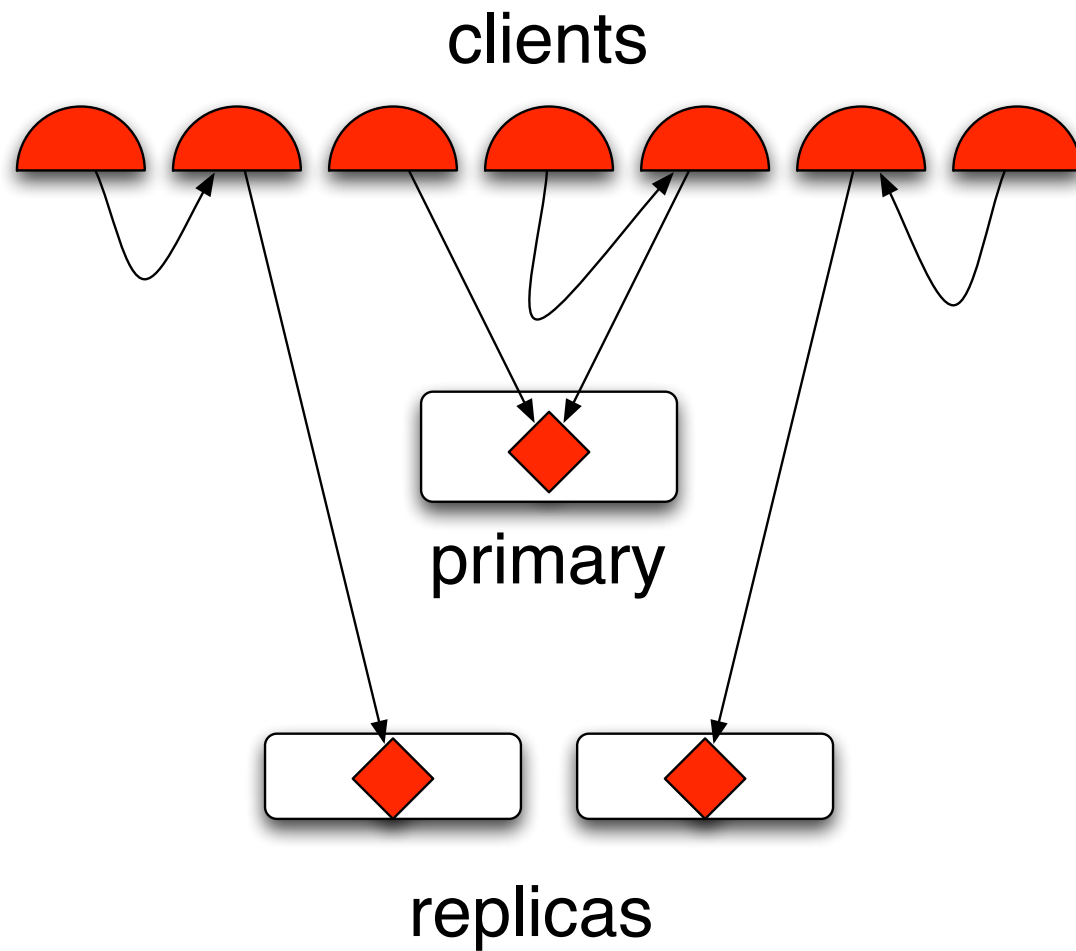
Read flash crowd




Read shedding




Read shedding



Write flash crowd



Write shedding




Write shedding

clients




Serialization:

Easy with POSIX
HPC I/O extensions,
complicated otherwise



primary



replicas

Preliminary Results

Throughput with overloaded OSD 0

