

Distributed Storage Systems

John Leach
john@brightbox.com
twitter @johnleach

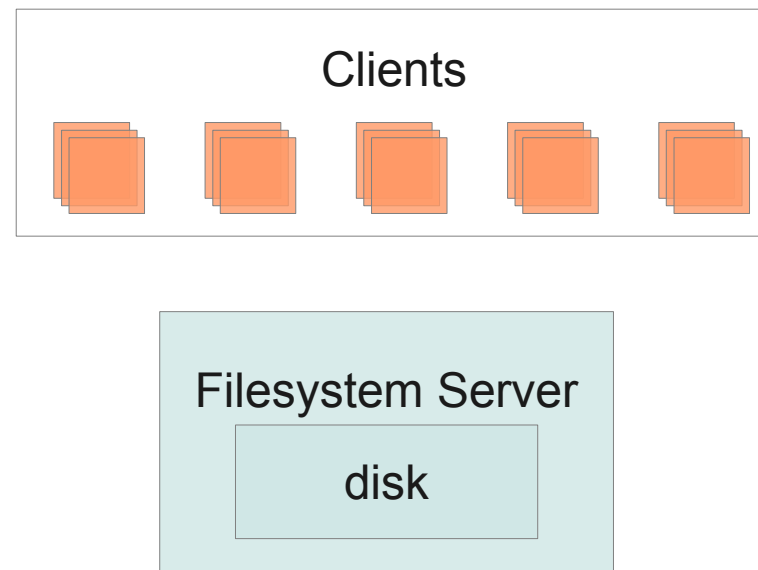
Brightbox Cloud
<http://brightbox.com>

Our requirements

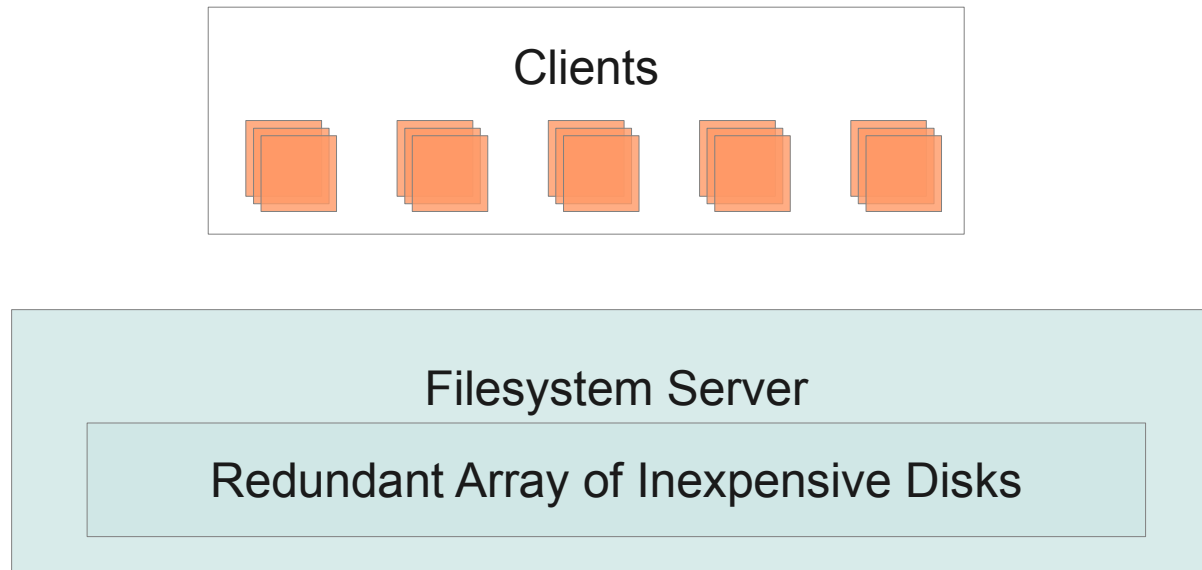
- Bright box has multiple zones (data centres)
- Should tolerate a zone failure
- Scale smoothly as data size grows
- Should use exciting unproven technology
- Libre software license

Brief history of file access



Scaling NFS: One disk



Scaling NFS: RAID

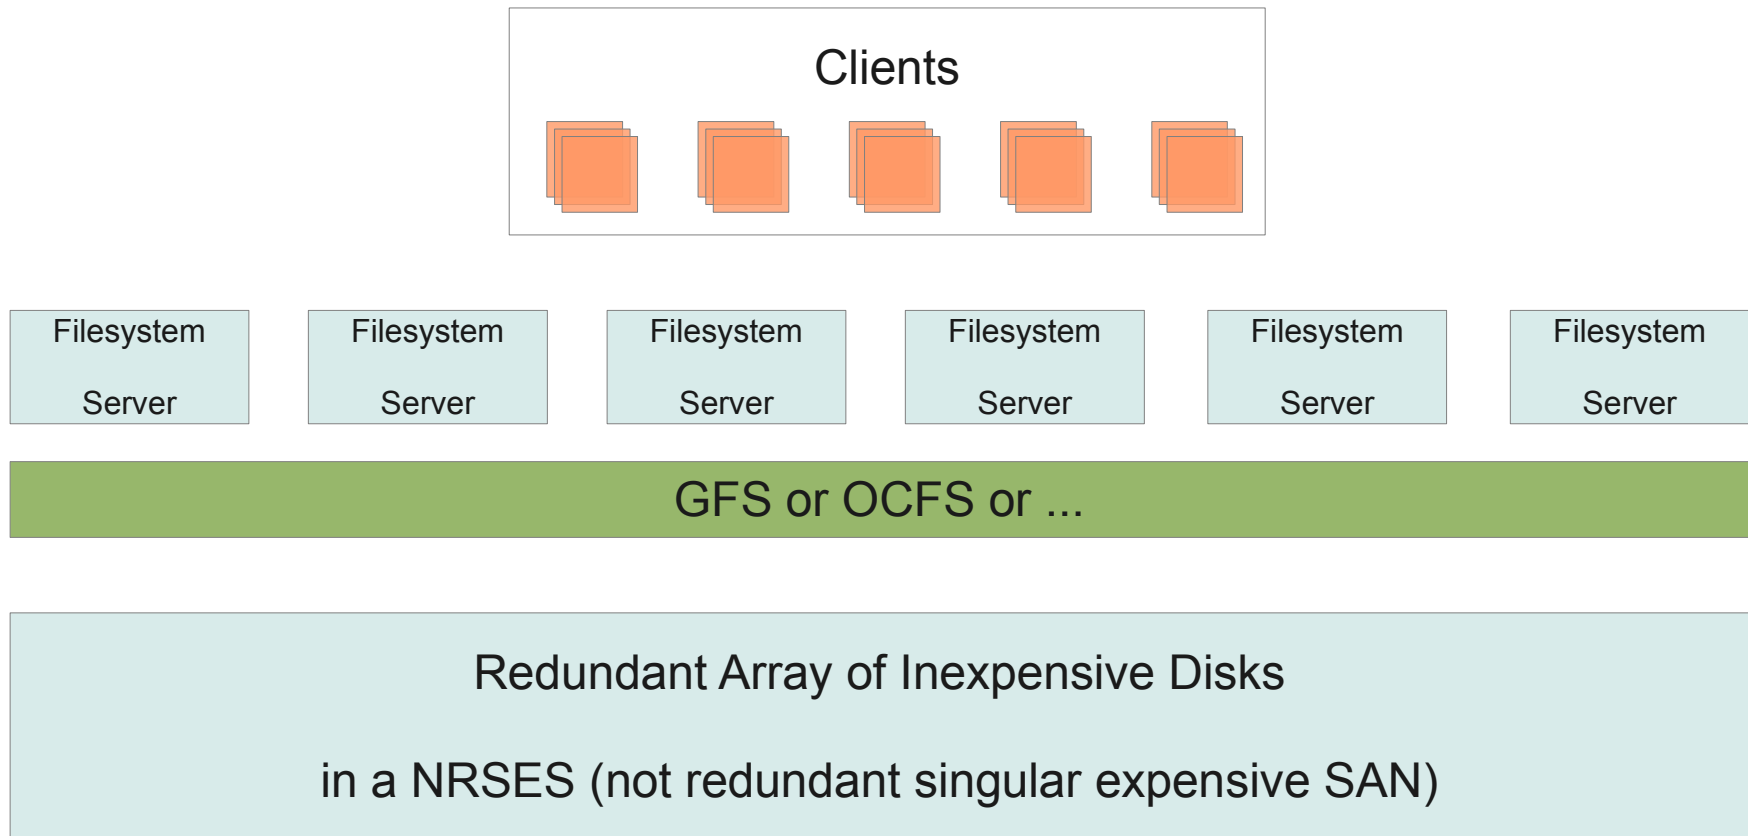


Scaling NFS: SAN

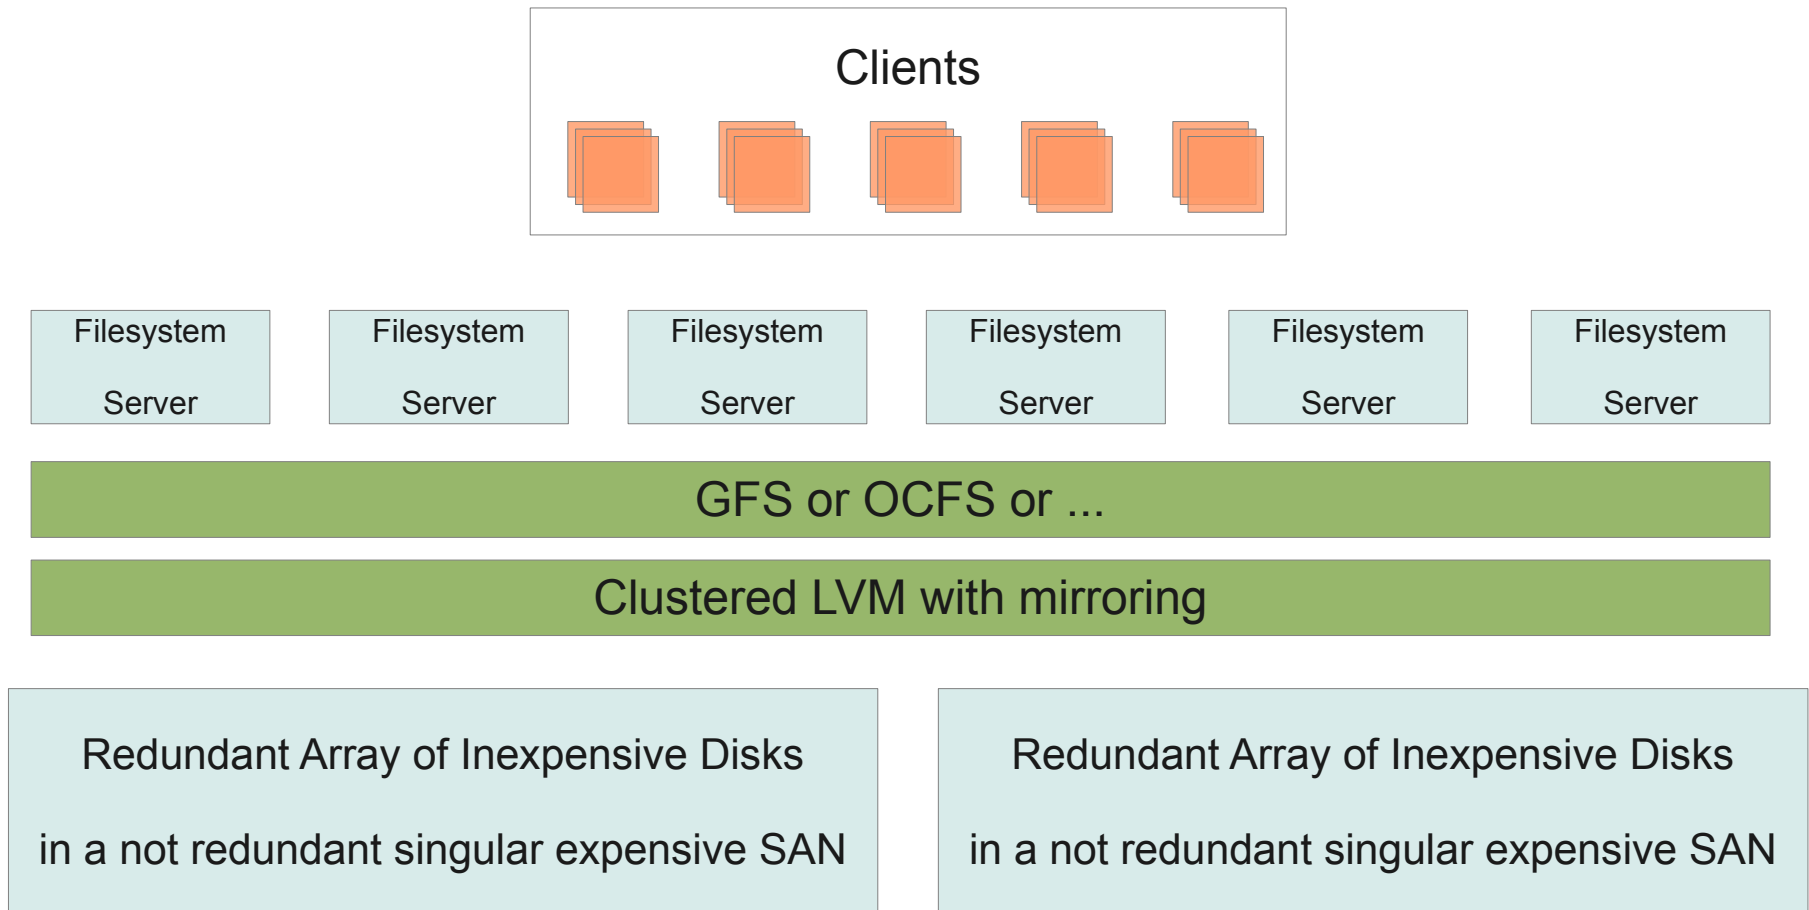


Redundant Array of Inexpensive Disks
in a NRSES (not redundant singular expensive SAN)

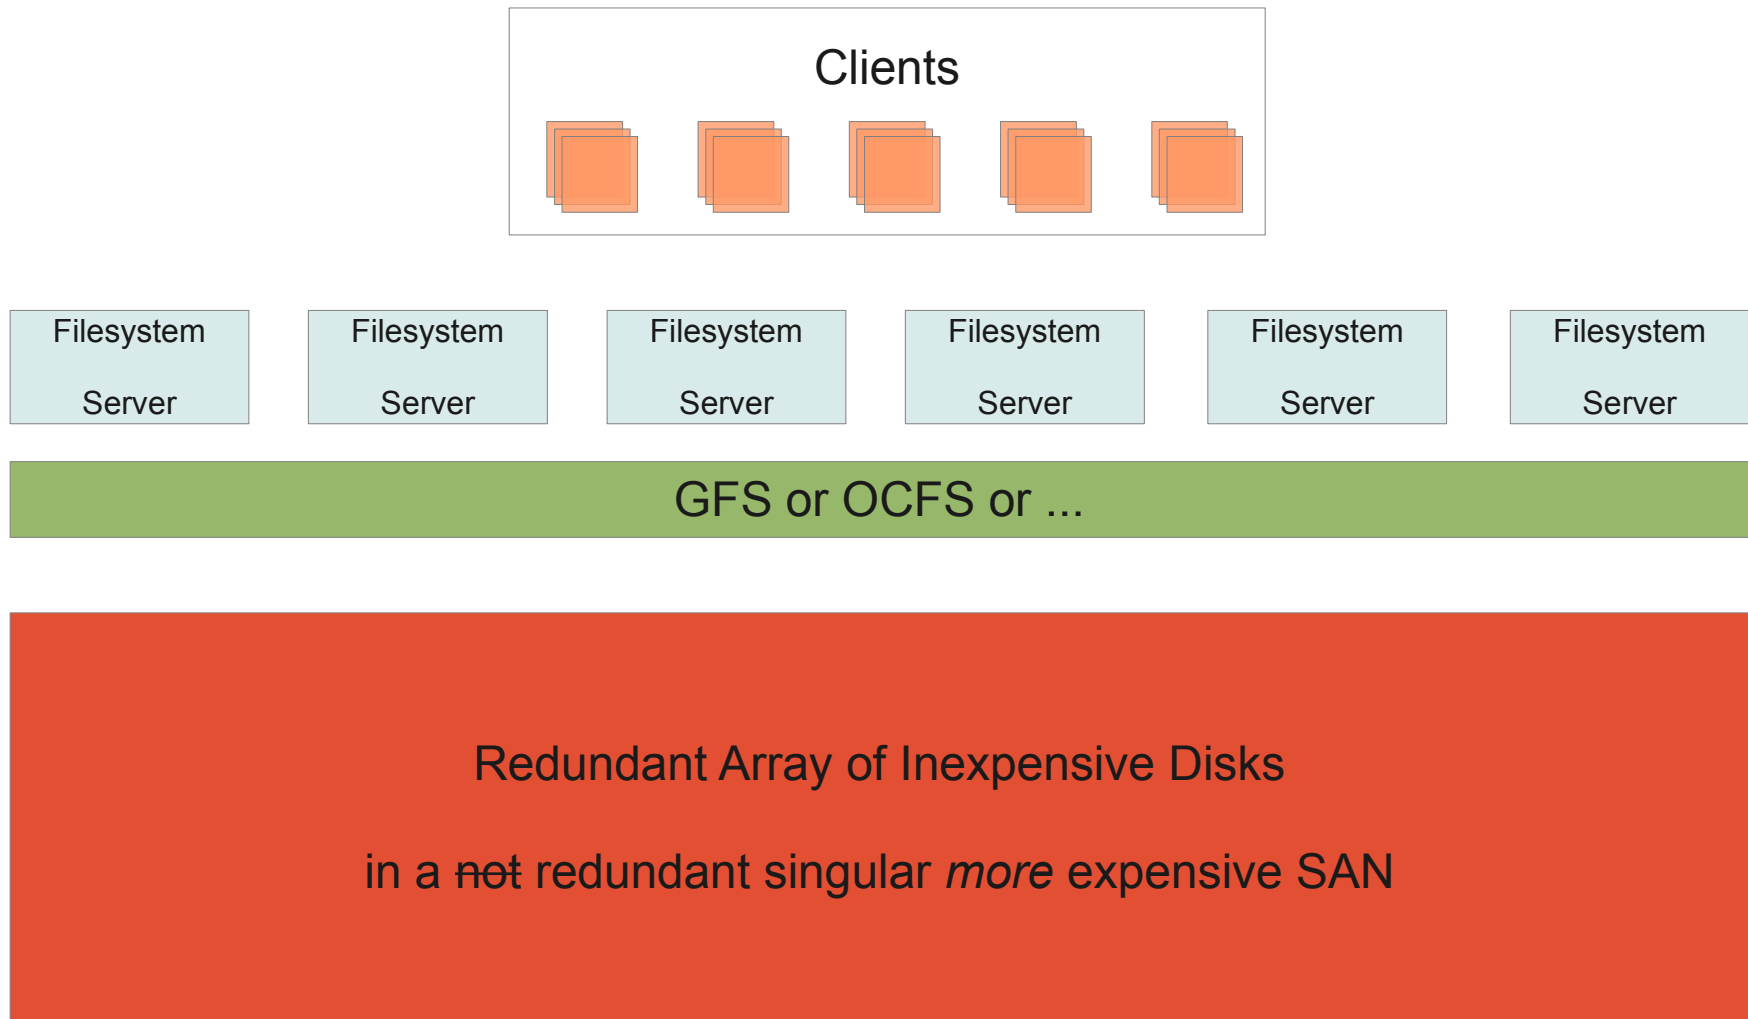
Scaling NFS: Shared disk fs



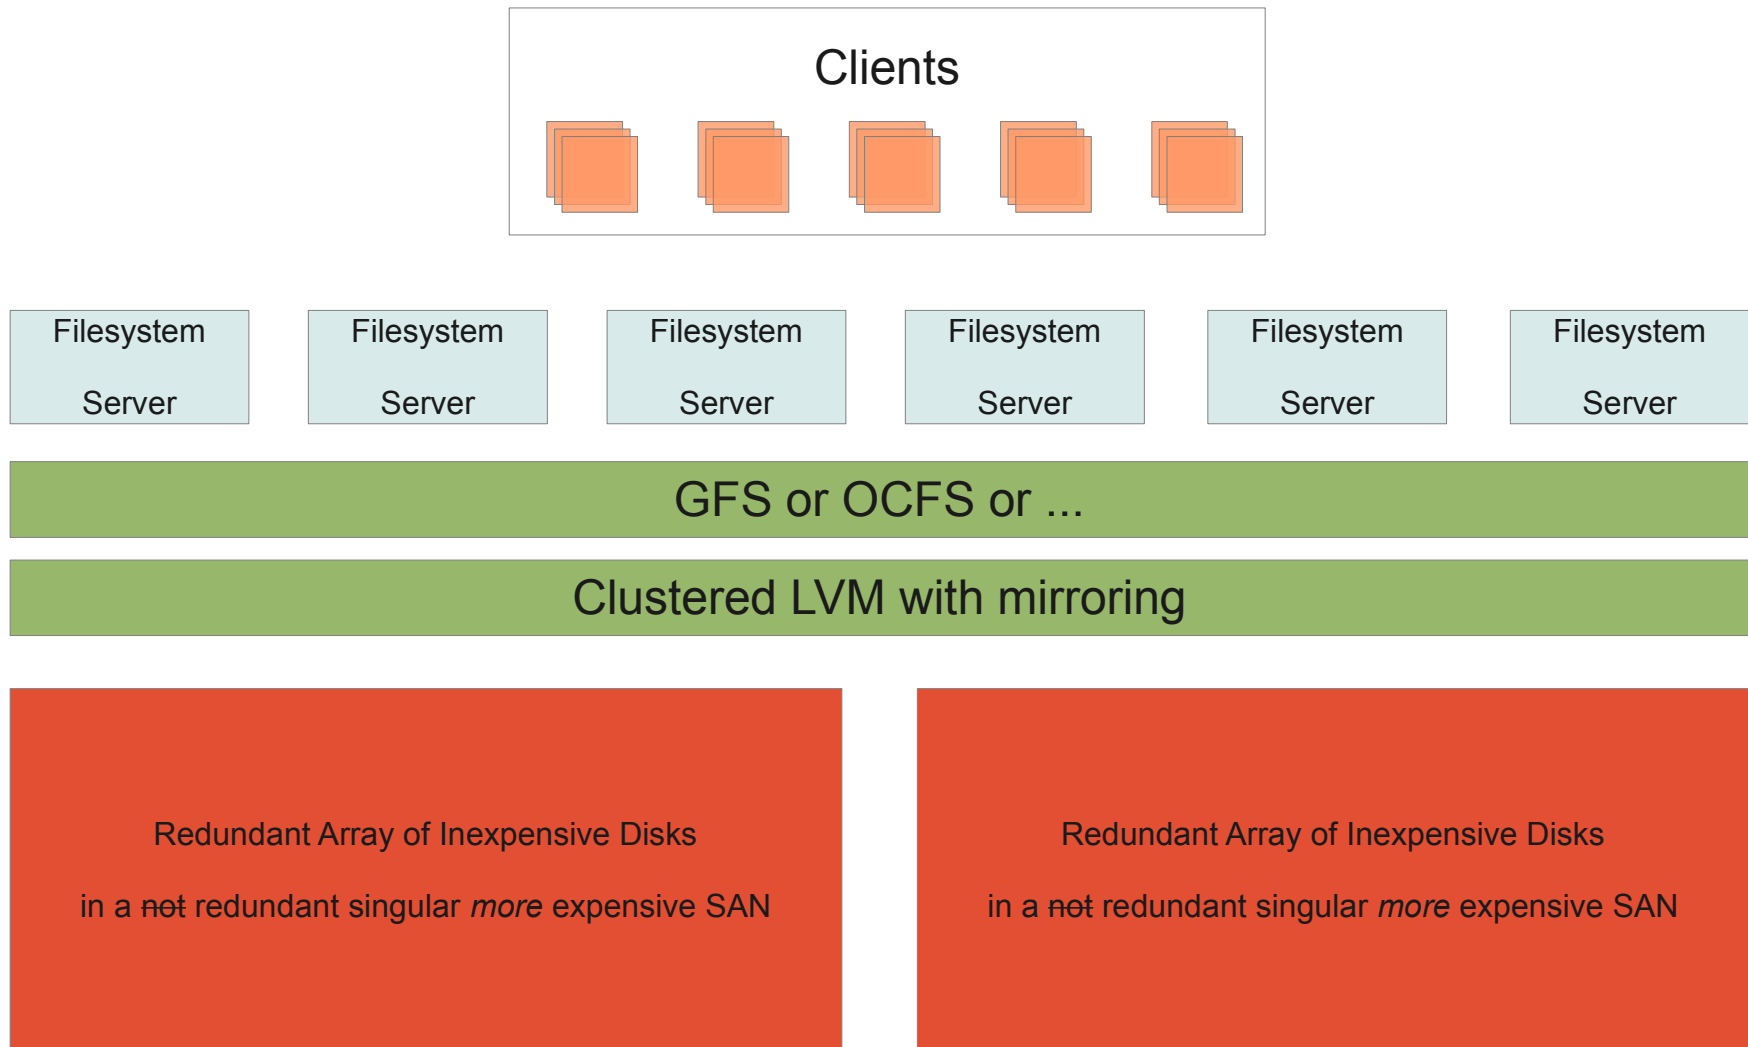
Shared disk fs: Replication



Shared disk fs: Replication



Shared disk fs: Replication



Old techniques

- Hot or warm standby servers
- Expensive SAN hardware
- Shared block devices
- Moving IP addresses
- Server side replication
- Scales mostly vertically
- Manual partitioning to scale horizontally

New techniques

- Shared nothing
- Clever clients
- Automatic partitioning
- Automatic replication
- Clever stuff: DHT, Vector clocks, PAXOS, Mapreduce, Merkle trees, Unicorn hooves
- ~~POSIX~~

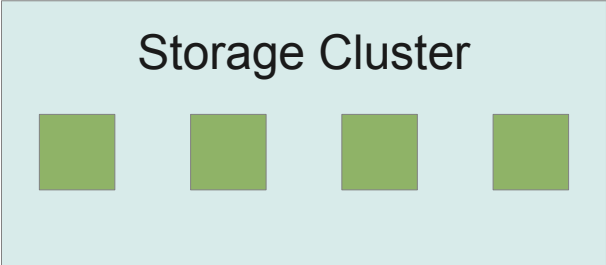

New Problems

- Locating your data
- Ensuring consistency
- Something has to give


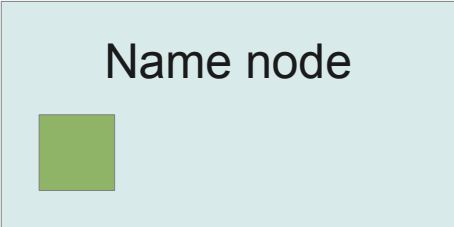

Brewers CAP theorem

- Consistency
- Availability
- Partition tolerance

GlusterFS



Hadoop File System



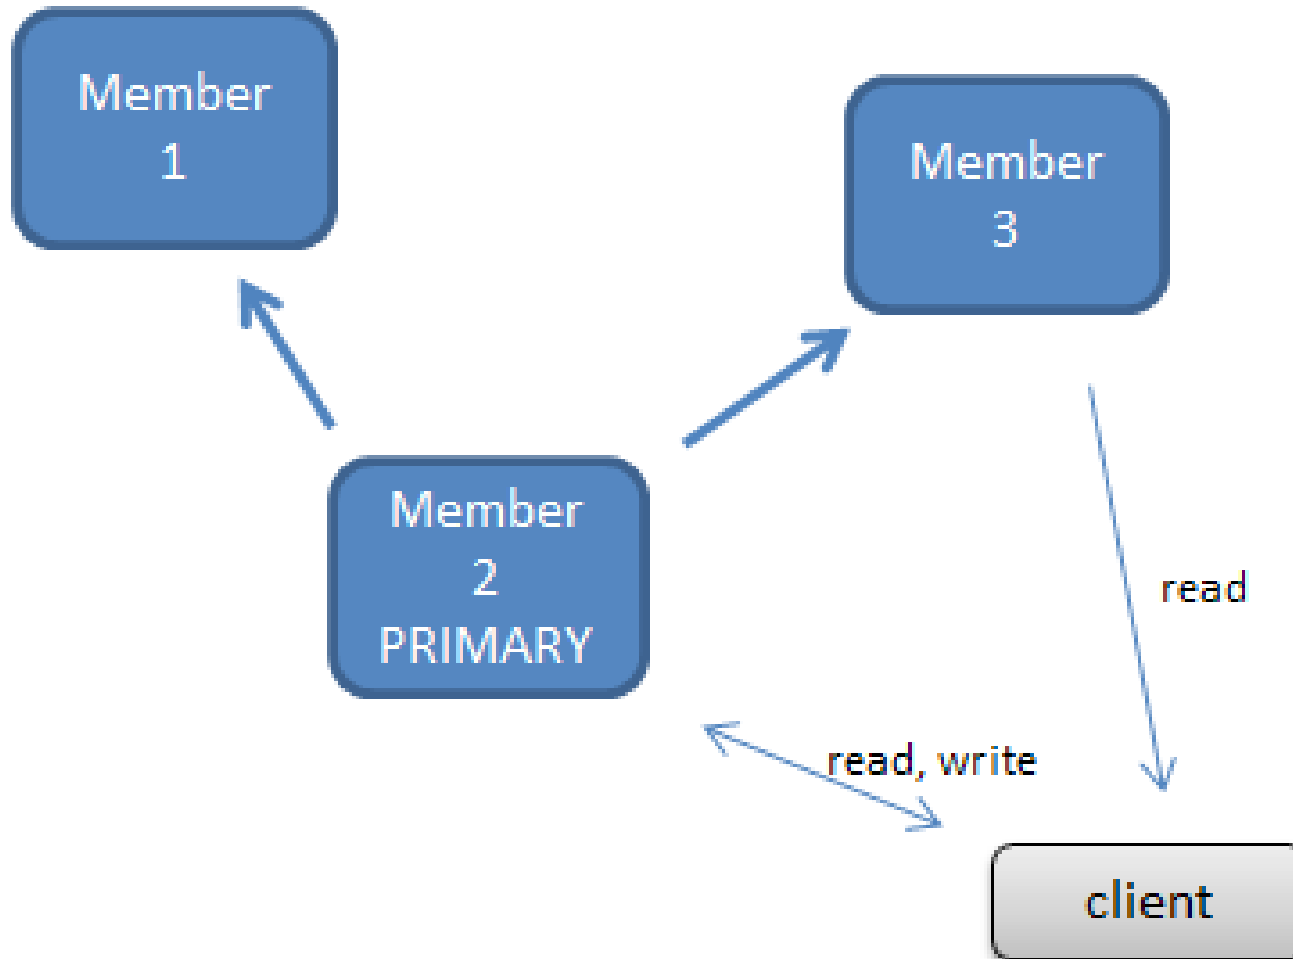
Hadoop File System

- Hot failover patches in Feb
- Batch processing, not interactive
- High throughput, not low latency
- Map Reduce
- Namenode SPOF
- Multi-data centre
- Consistent



MongoDB

- Document store, dynamic schema
- Async replication
- Primary server for writes
- Automatic sharding
- Map Reduce
- GridFS for large files
- Multi-datacentre, but not partition tolerant
- Mostly consistent

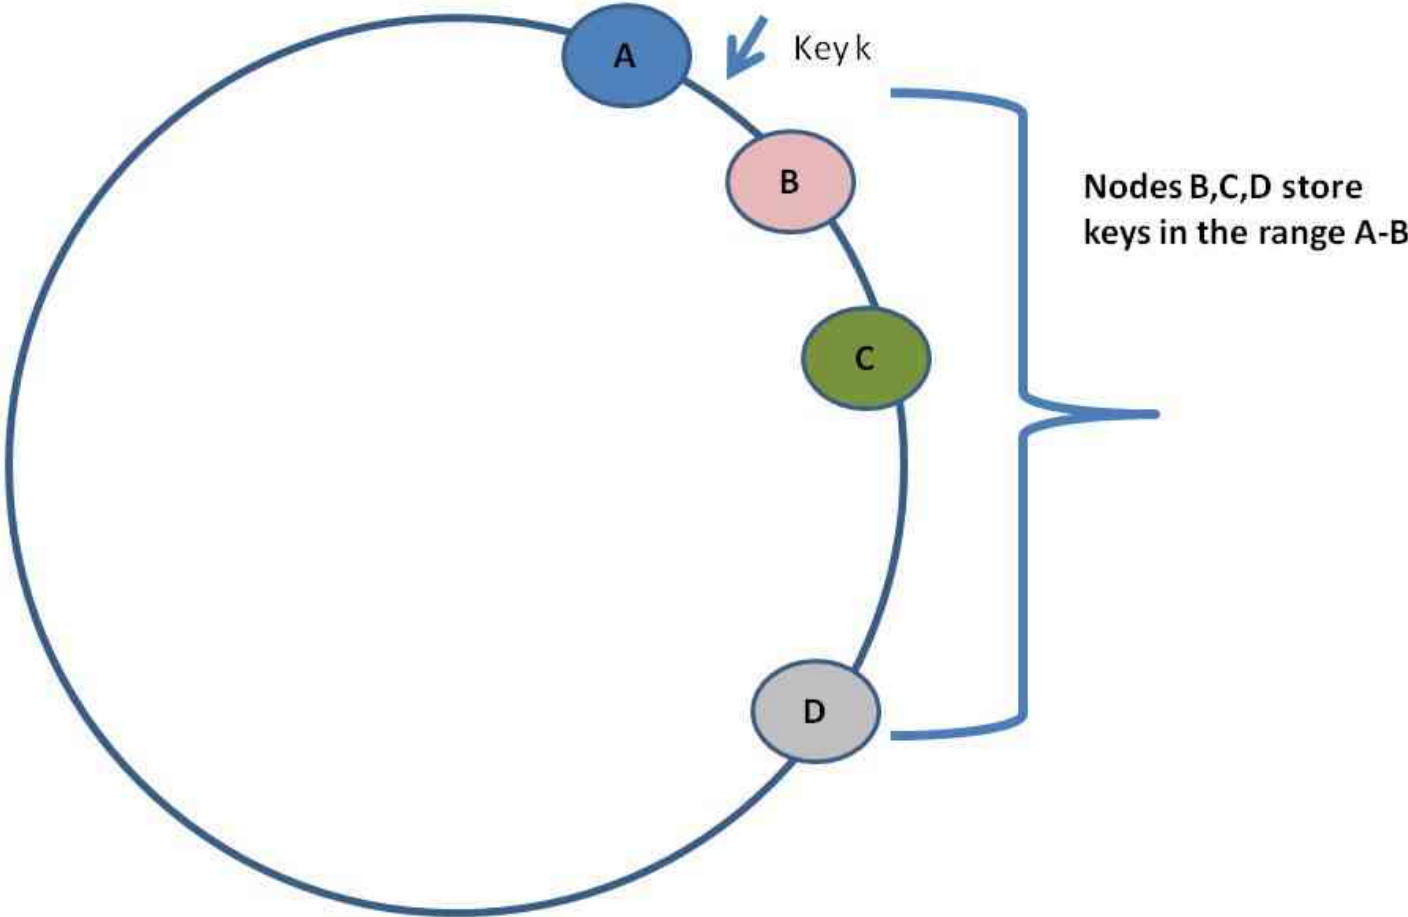
MongoDB



Openstack Swift



Openstack Swift



Cassandra

- P2P, DHT, Gossip, Hinted Handoff
- Column orientated. Data ordered.
- Design schema for types of queries
- Very fast highly available writing
- Per request consistency. Multi-data centre
- Thrift API

Riak

- Key value store.
- DHT, Gossip, Vector Clocks
- Map reduce
- Luwak for large files





Zookeeper

- PAXOS like consensus protocol
- Read scales up with more servers
- Writes slow down with more servers
- Always consistent
- In-memory
- Strict ordering
- Small data

Ceph

- Object store
- Full POSIX file system on top
- PAXOS for cluster state
- CRUSH rather than DHT
- Multi-datacenter.
- Strongly consistent, not partition tolerant
- RBD, S3-alike, plus POSIX

Ceph



Distributed Storage Systems

John Leach
john@brightbox.com
twitter @johnleach

Brightbox Cloud
<http://brightbox.com>