

Graduate Operating Systems

Spring 2022

1

Paper “Ceph”

- *What are the goals of Ceph?*
- Object Storage Devices
- Metadata Servers
- *Why does OSD suffer from scalability problems?*
- POSIX
- Consistency & coherence
- Design feature 1: Decoupled data & metadata
- Design feature 2: Dyn. distr. metadata mgt.
- Design feature 3: Reliable autonomic distr. object storage

2

Paper “Ceph”

- Capabilities
- Passive/active “live messages”

3

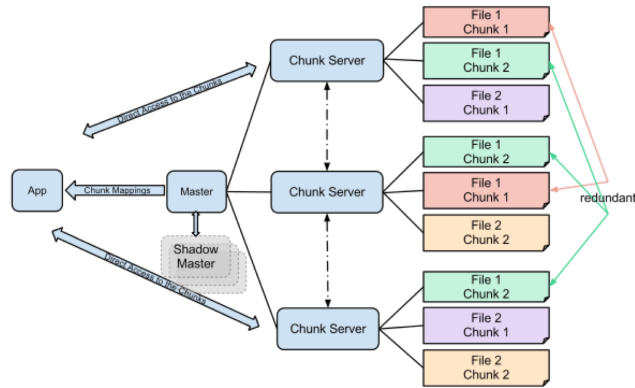
Paper “Google FS”

- Assumptions:
 - Component failures are the norm, not the exception
 - Files are huge by traditional standards
 - Most file updates are append-only
 - *How does this compare to previous papers?*
- System is built from many inexpensive commodity components
- System will store a modest number of large files

4

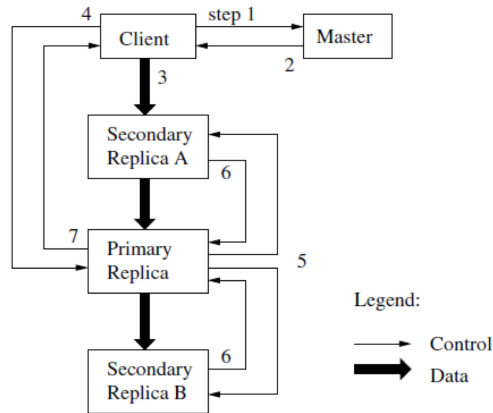
Paper "Google FS"

- Non-POSIX, "snapshot", "record append"



5

Paper "Google FS"



6

Paper “Google FS”

- Single master
- Leases, heartbeat messages, namespace management
- Re-replication, rebalancing, garbage collection

- *How does it compare to Ceph?*
- *Why does Google FS avoid file caching?*
- *What are the pros/cons of large chunk sizes?*

- Google “Colossus” (2010)
- “Single point of failure”