

Building a Scalable NFS

Breaking the Single Server Bottleneck

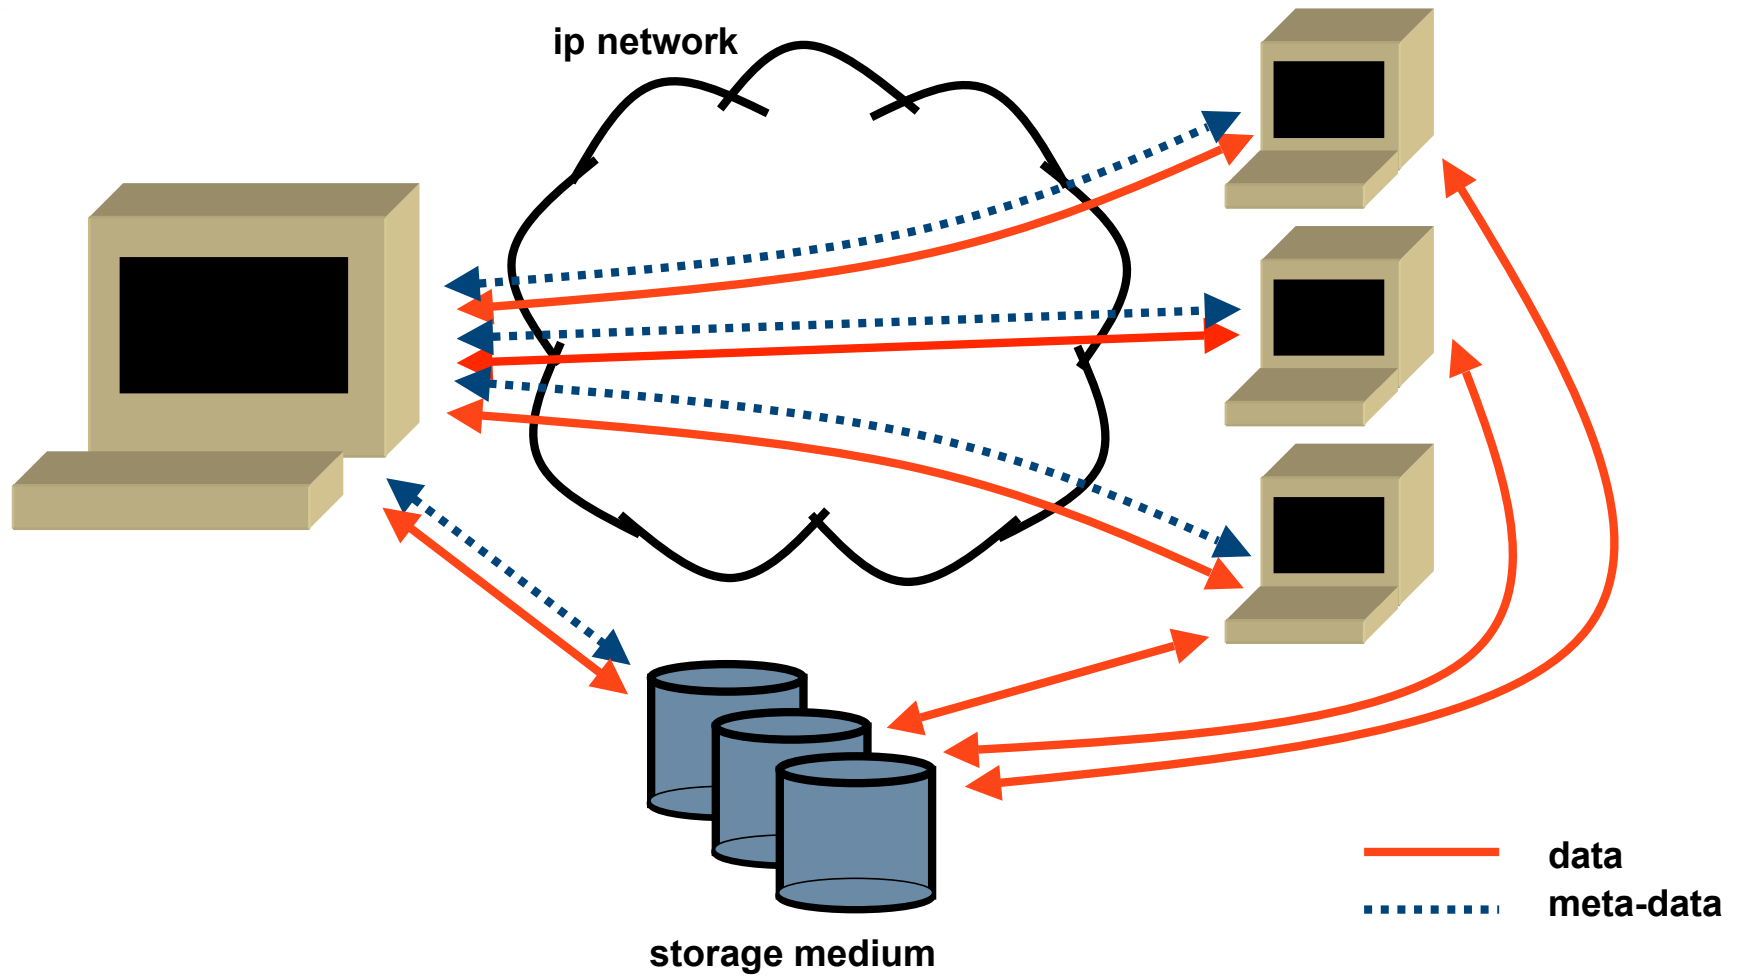
Stephen Fridella, Xiaoye Jiang, David Black

Goals

Extend NFS in a way which:

- Breaks the single server bottleneck
 - Supports multiple underlying storage models
 - Achieves scalable sharing for file systems *and* individual files
-


Single Server Bottleneck



What Storage Medium?

All of them...

File System




"logical volumes" ---- block containers


abstraction

realization


SCSI/ FC device



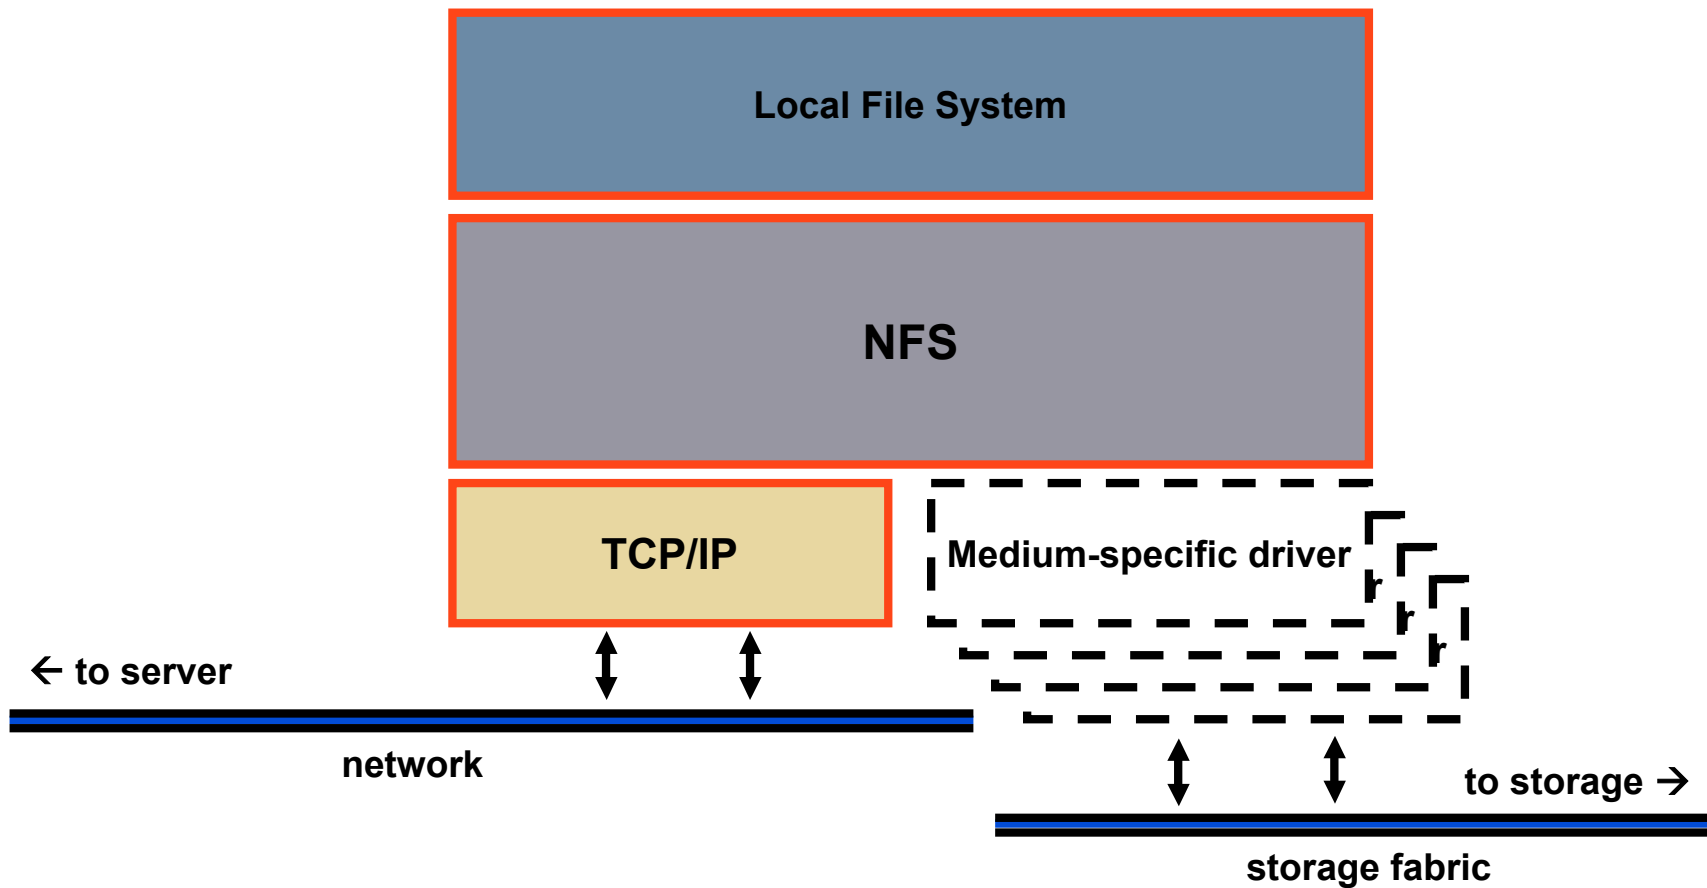
OSD



NFS/ Samba



NFS Client Architecture



Volume Discovery

Need volume discovery phase

- Client and server agree on what storage medium is used
 - Server communicates to client a mapping of logical volume ids to appropriate storage space
 - Future communications are in terms of logical volume addresses
-

Meta-Data Management Protocol

Client needs protocol elements which allow it to...

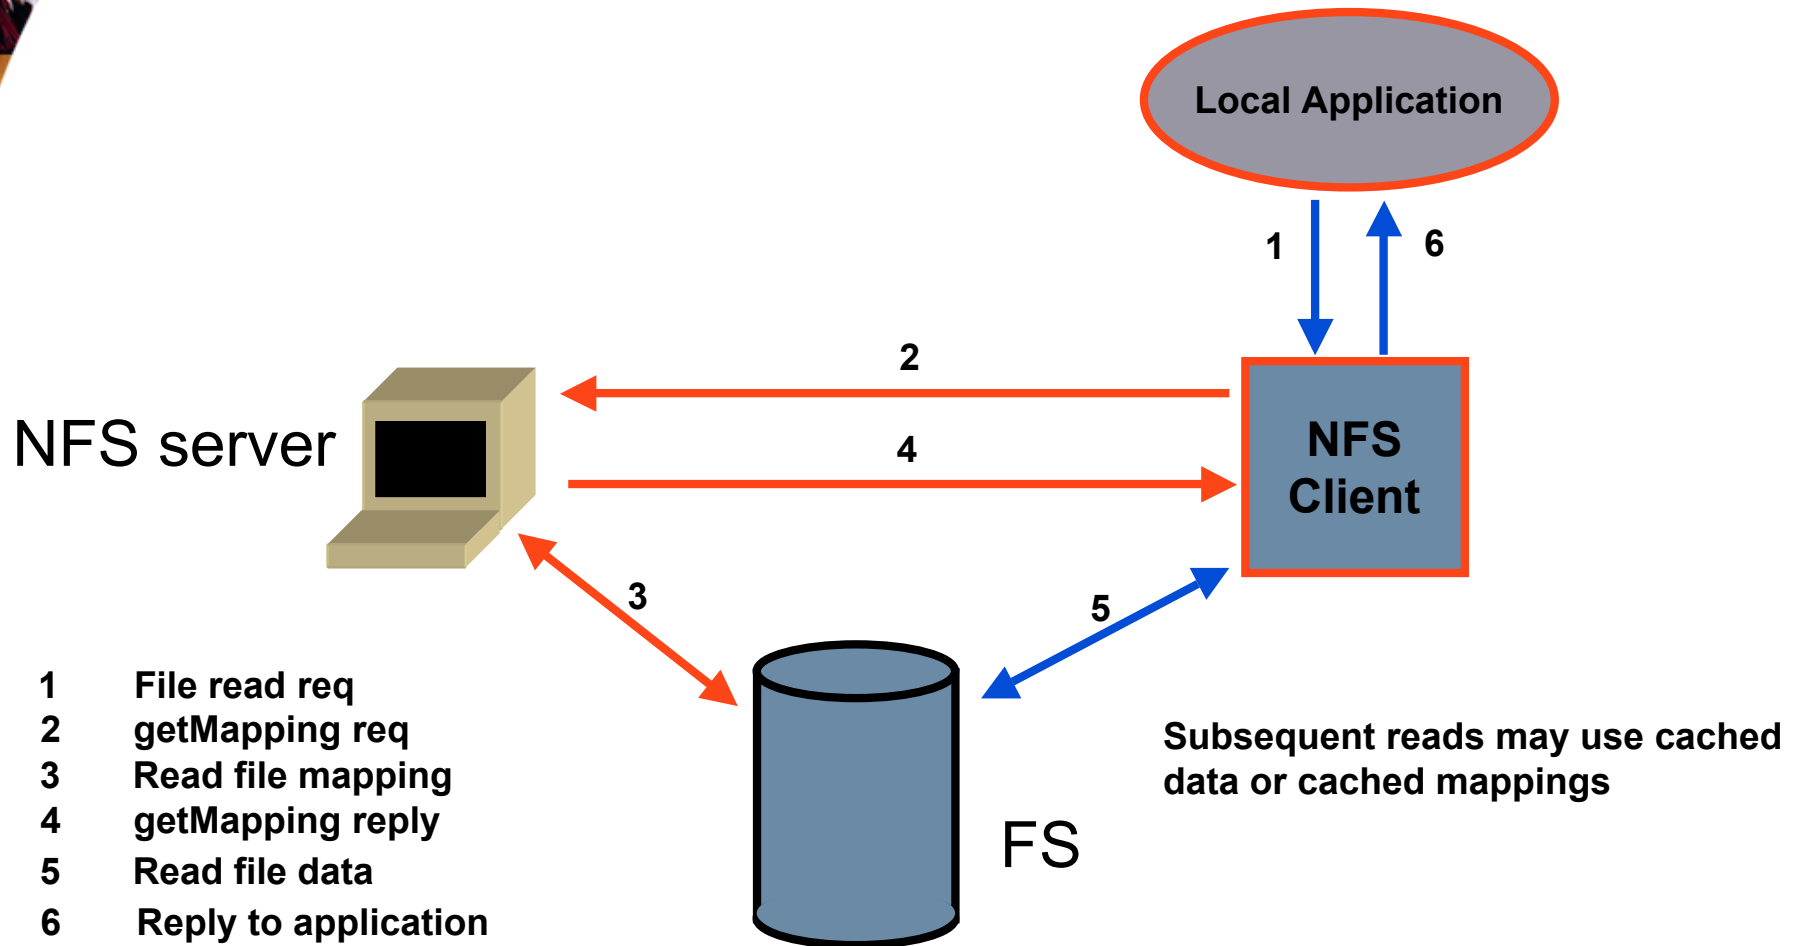
- fetch block mappings for files
 - allocate new blocks to files
 - coordinate access to files with other clients
-

Meta-Data Management (cont)

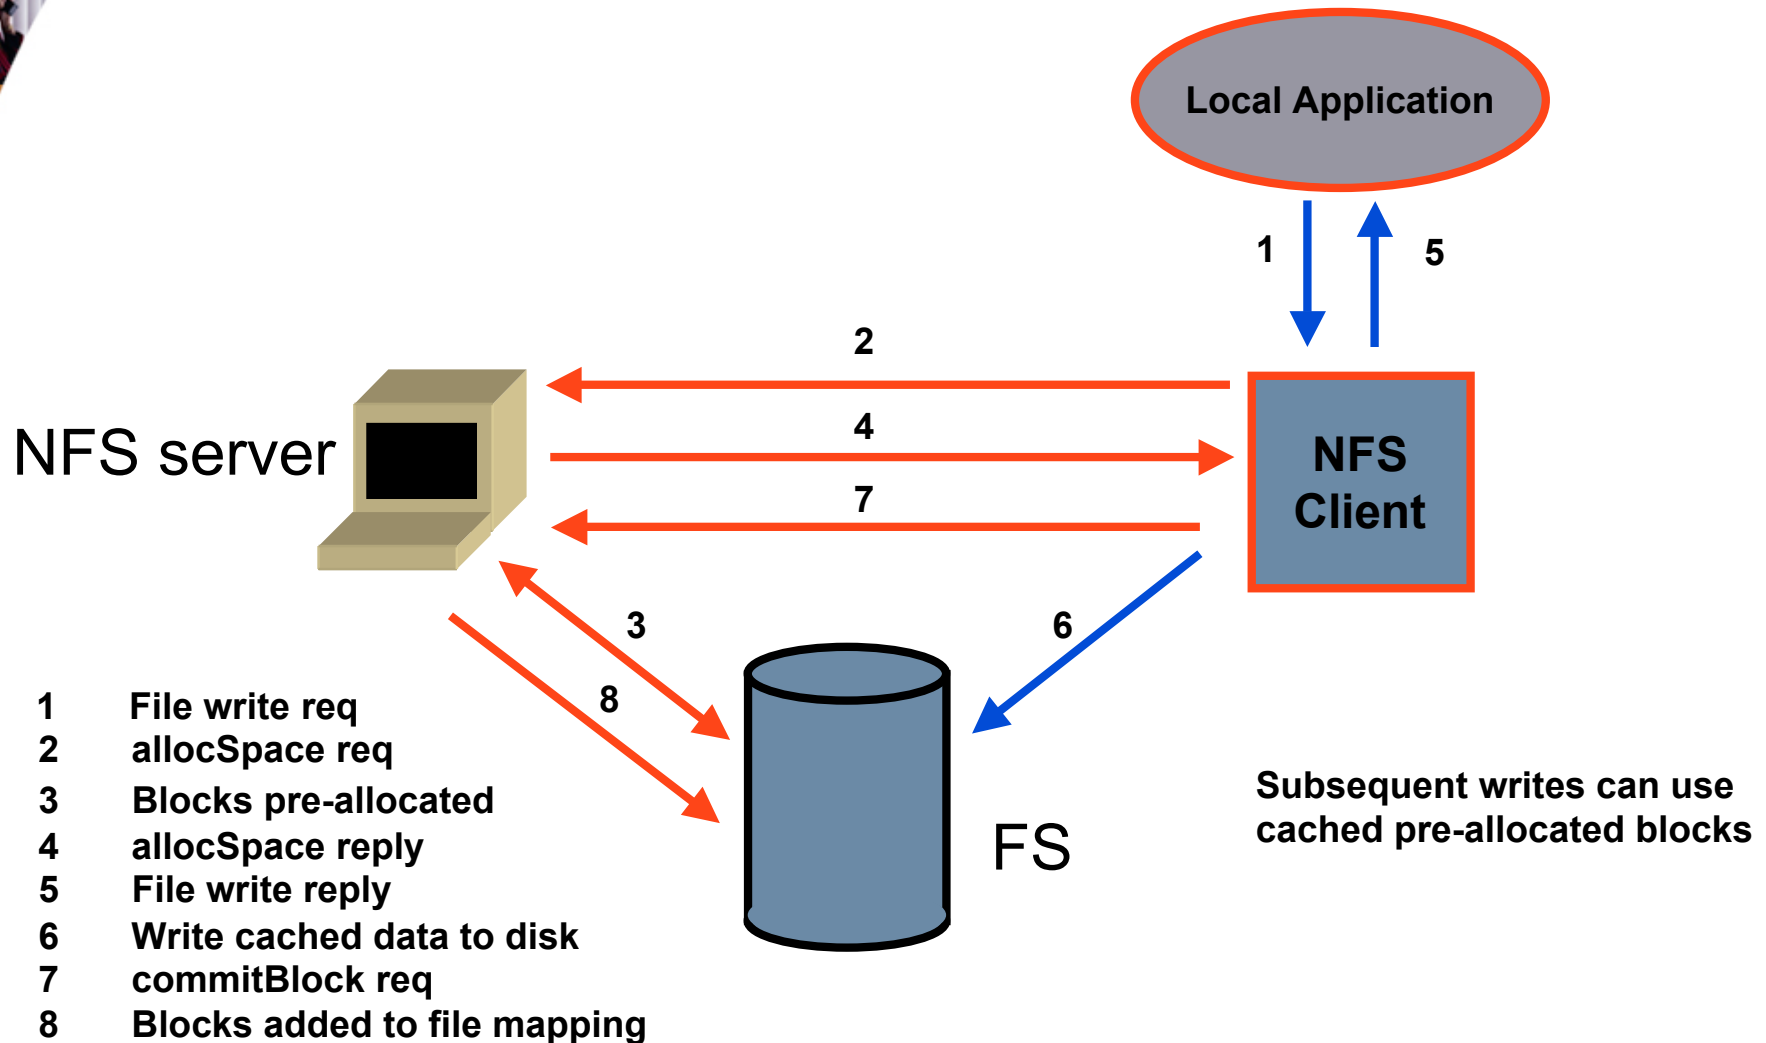
Our proposal:

- Block-range read and write delegations are implicitly granted on successful mapping/allocation requests
 - Allocations are initially provisional---they require an explicit commit request after data is written
 - Ideas are based on EMC's File Mapping Protocol
-

Reading Data (example)



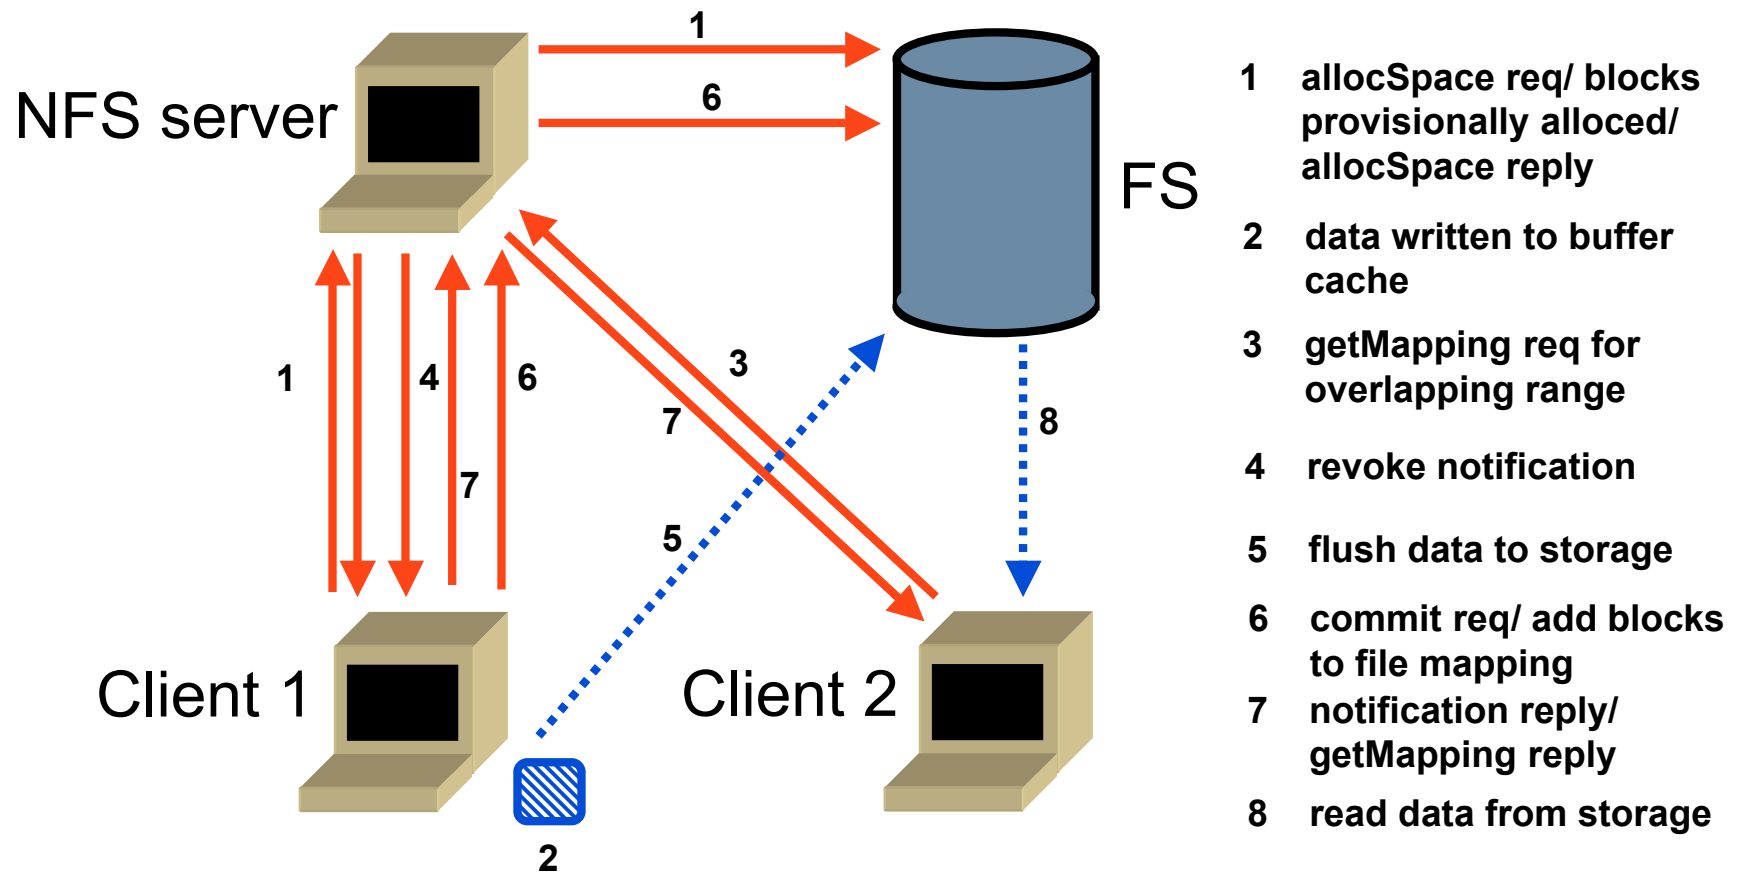
Writing Data (example)



Block Range Delegations

- Shared and exclusive
 - Lease-based
 - Can be revoked when server detects conflicting delegation request
 - Can be revoked when server detects operation that will change mapping (write, truncate)
-

Conflicting Requests (Example)



Summary

- Break the single server bottleneck by providing a meta-data management interface for clients
 - Support multiple underlying storage models by using a generic logical volume abstraction
 - Provide for read/write sharing of files as well as file-systems by using block-range delegations
-