



HECURA Research Program at Clemson

Walt Ligon
Clemson University

Program Areas Addressed

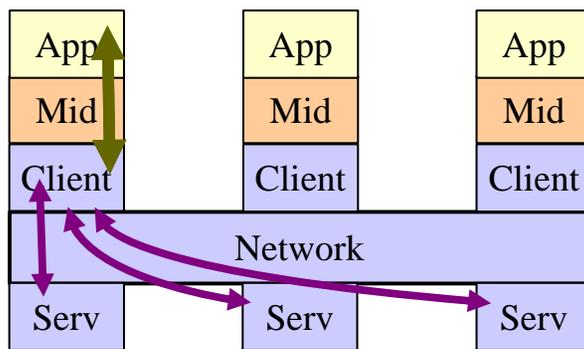
- scalable metadata operations
- scalable small and unaligned operations
- I/O middleware
- active caching
- server to server communication
- simulation of I/O, file, and storage systems

Research Focus

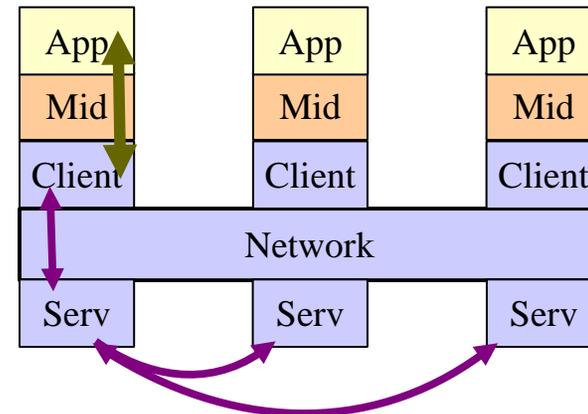
- Server-to-server communication
 - scalable metadata
- Run-time configurable semantics
 - Lockless SC semantics
 - POSIX versus weakened models
 - caching

Scalable Metadata

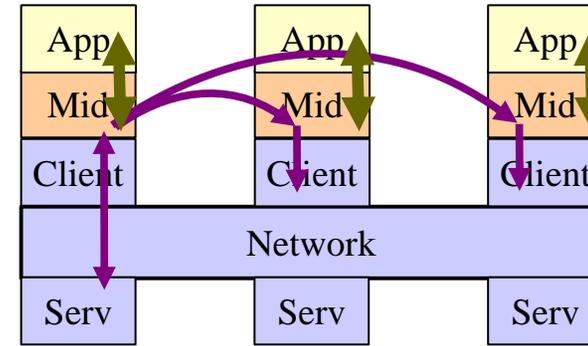
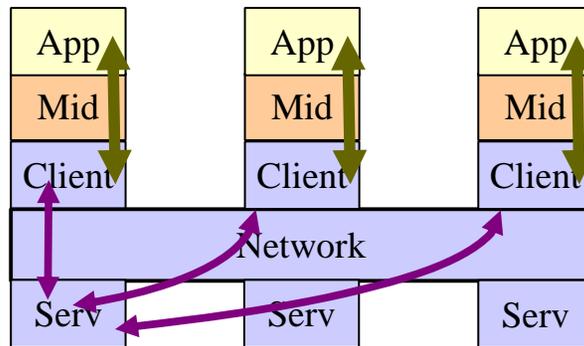
Server-to-Server Communication



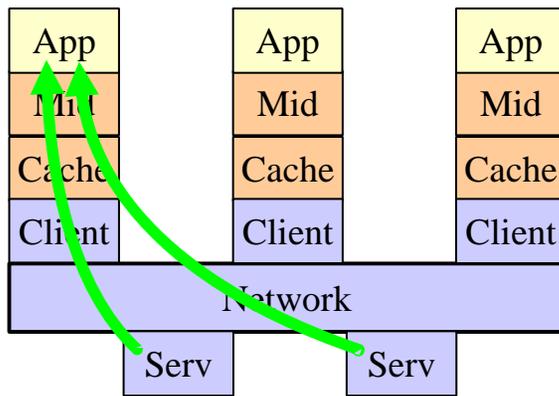
Traditional Metadata Operation



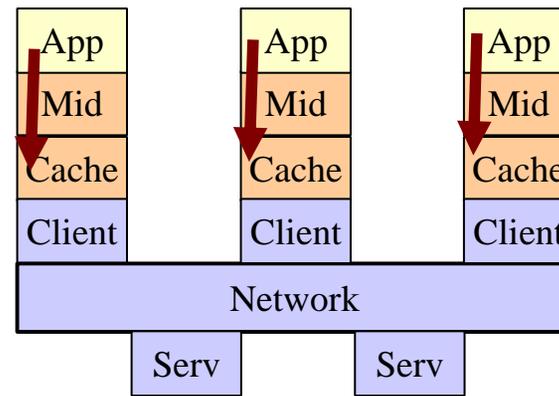
Scalable Metadata Operation



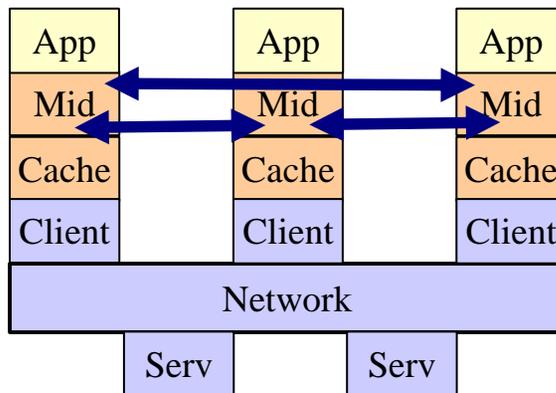
Middleware Managed Cache Weakened Consistency



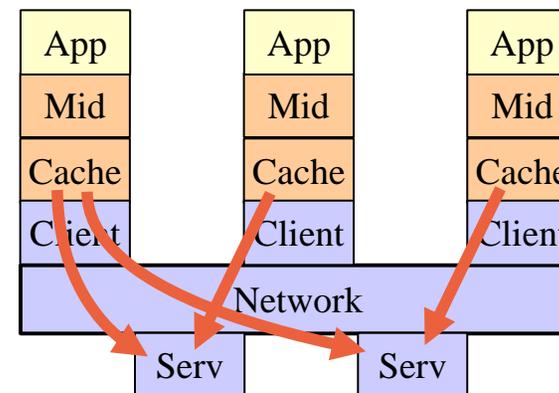
Cold Read



Write



Synchronization
Event

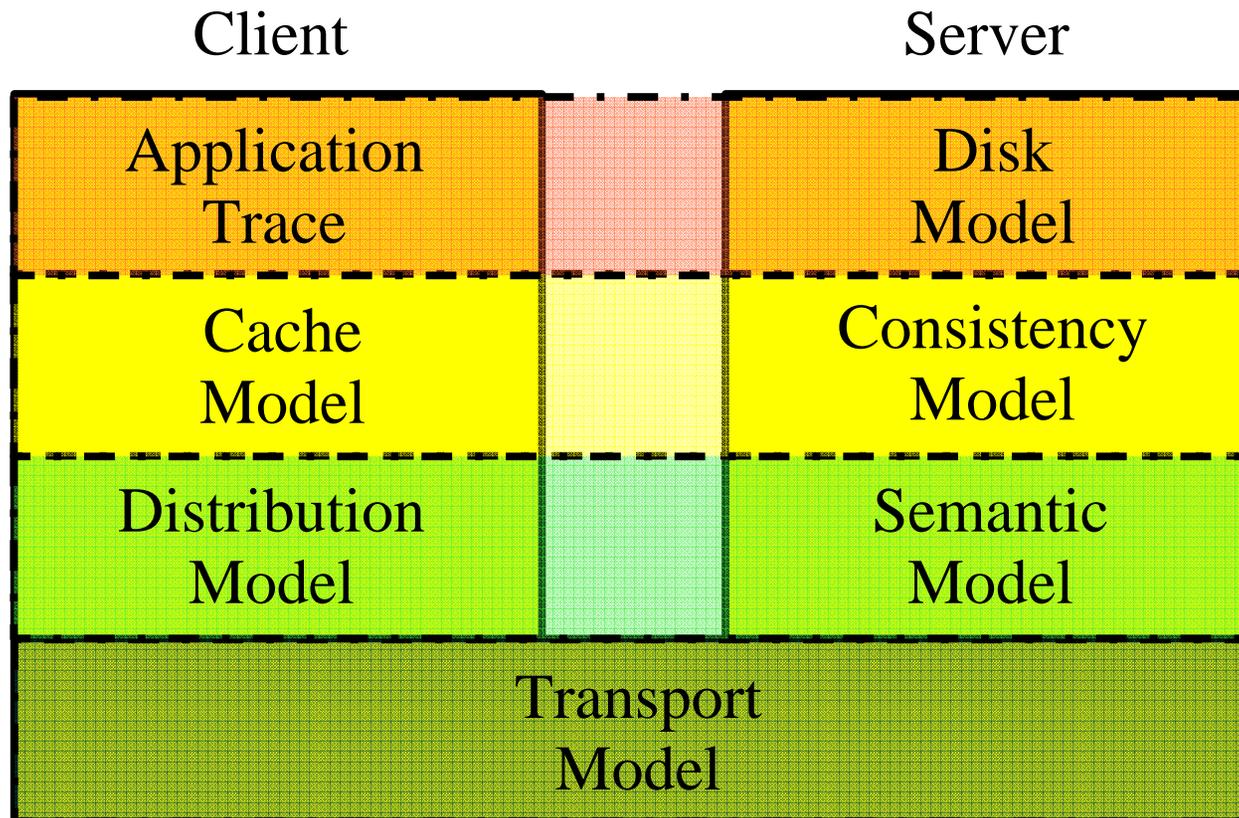


Write-back

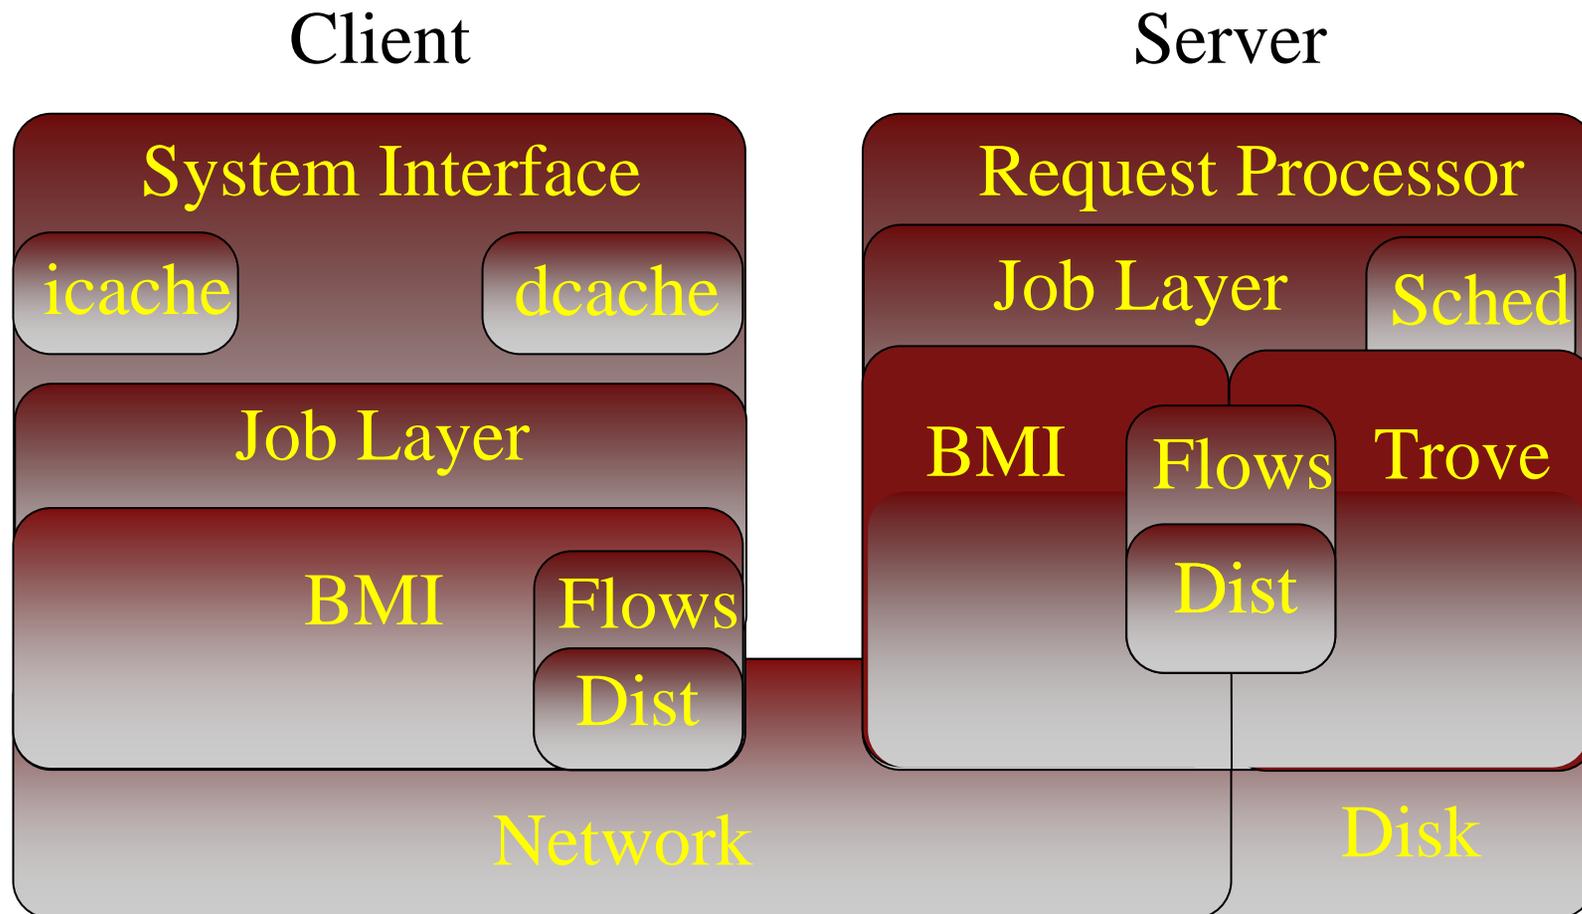
Research Approach

- File system simulation environment
 - Extensible
 - Based on PVFS2
- Middleware implementation

File System Simulation



PVFS2 Architecture Simulation Model



I could go on for hours ...

PVFS 2

PARALLEL VIRTUAL FILE SYSTEM