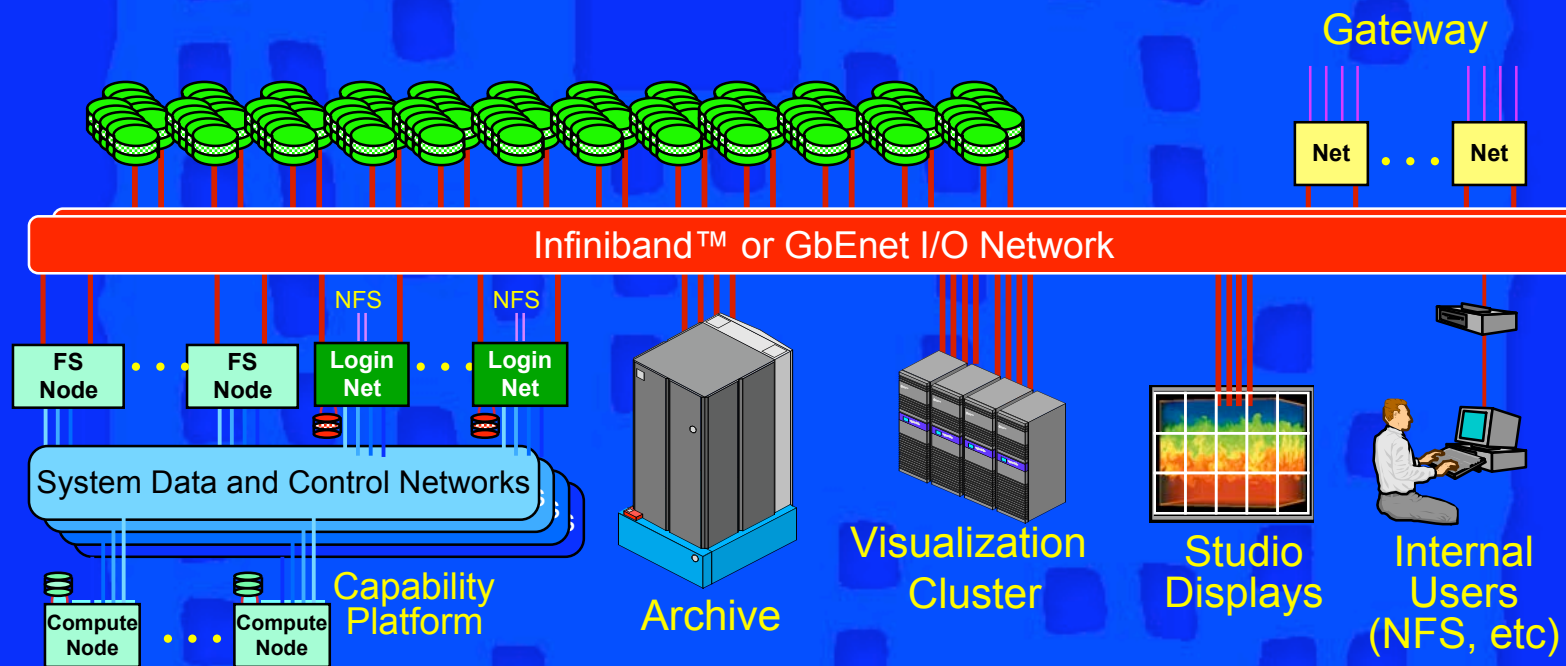


Enabling Parallelism in NFSv4

* Peter Honeyman
Center for Information Technology Integration
University of Michigan, Ann Arbor

Storage complexity

ASCI Platform, Data Storage and File System Architecture



from ASCI Technology Prospectus, July 2001



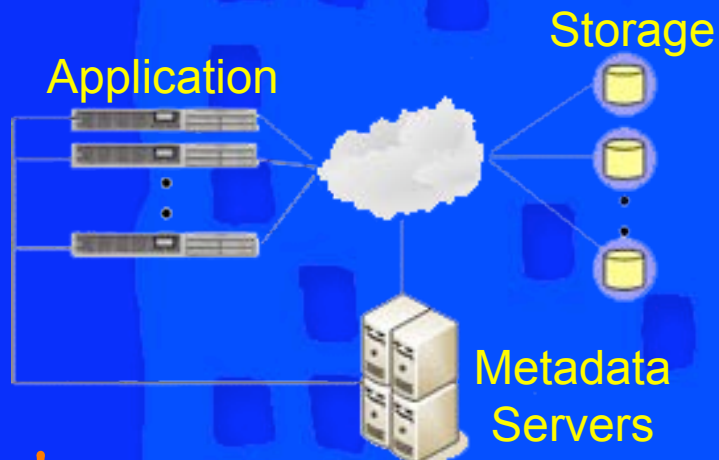
Parallel file systems

◆ Asymmetric

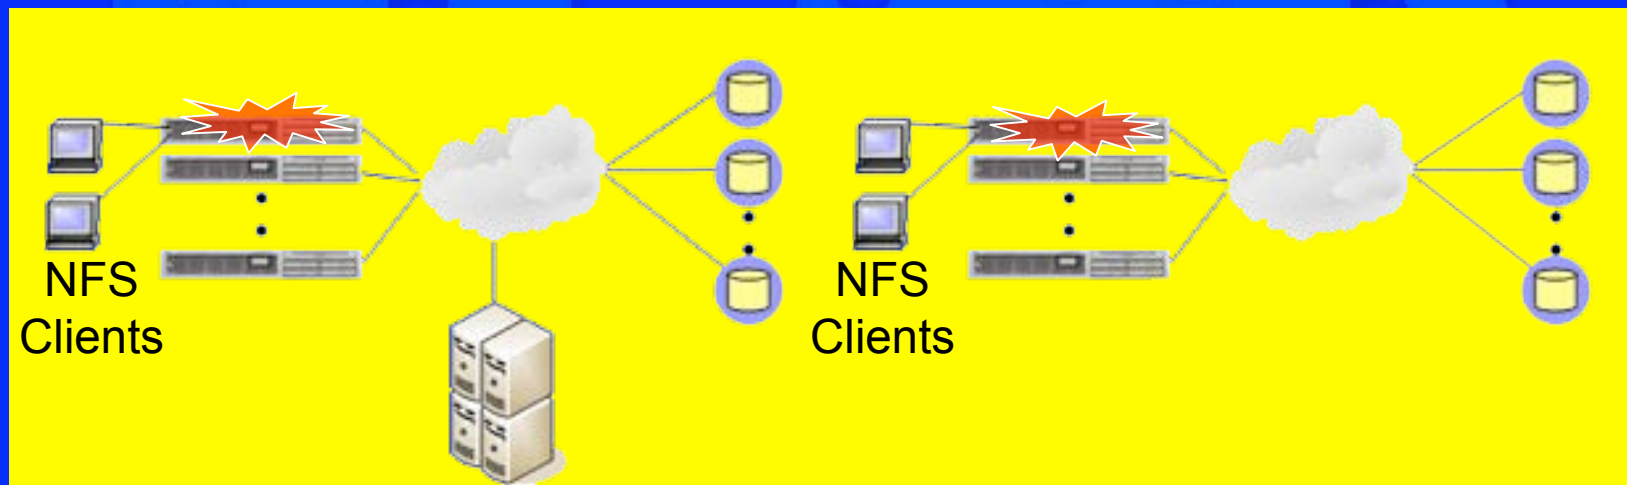
- ◆ Direct access to storage
- ◆ Separate metadata server(s)
- ◆ File based (Expand, NASD NFS, PVFS2)
- ◆ Object based (Lustre, Panasas, ActiveScale)
- ◆ Block based (EMC High Road, IBM SAN FS)

◆ Symmetric

- ◆ Direct access to storage
- ◆ Each node is a fully capable client and metadata server
- ◆ Block based (IBM GPFS, Redhat GFS, Polyserve Matrix Server)



NFS advantages and obstacles



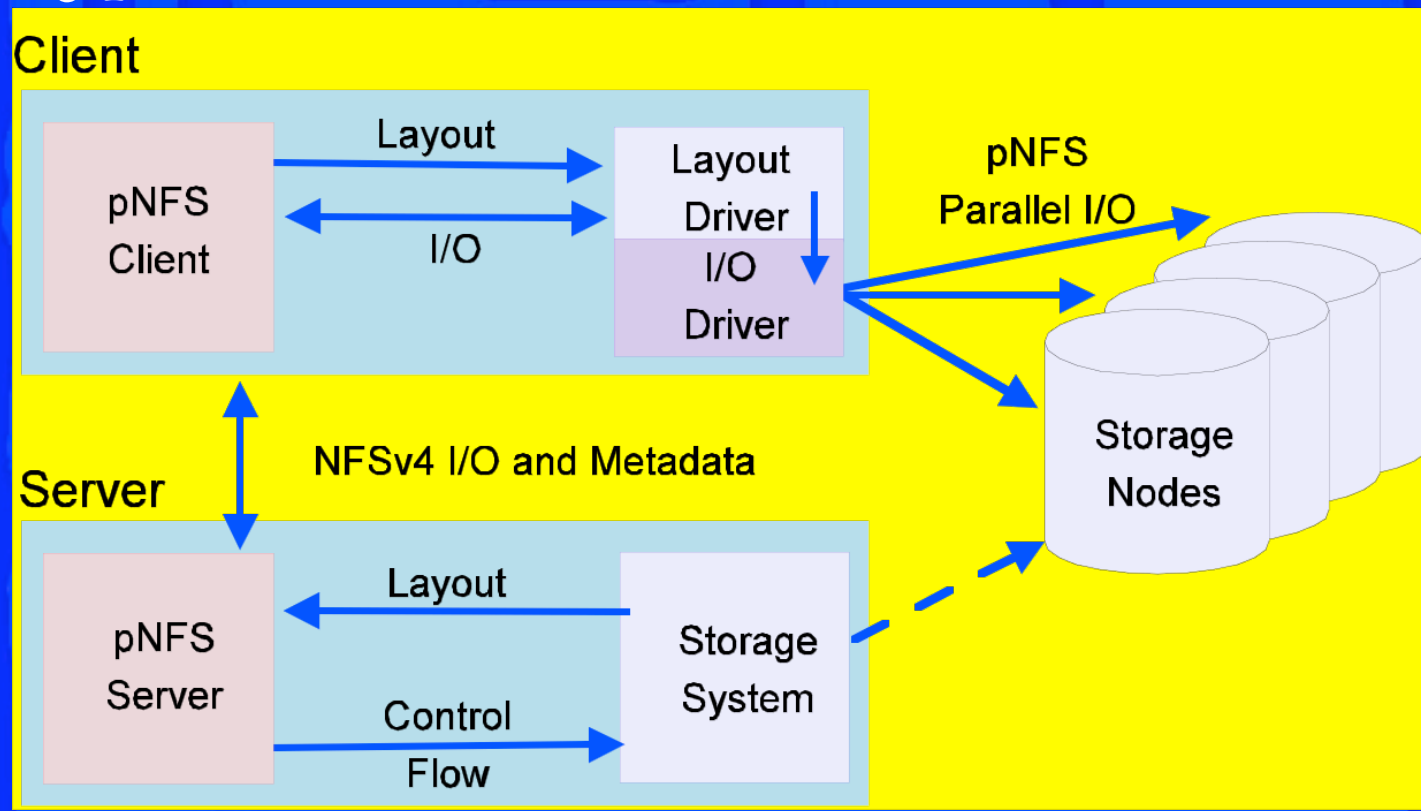
- ✓ Security
- ✓ Heterogeneity
- ✓ Transparency

- ✗ Poor performance
- ✗ Not scalable

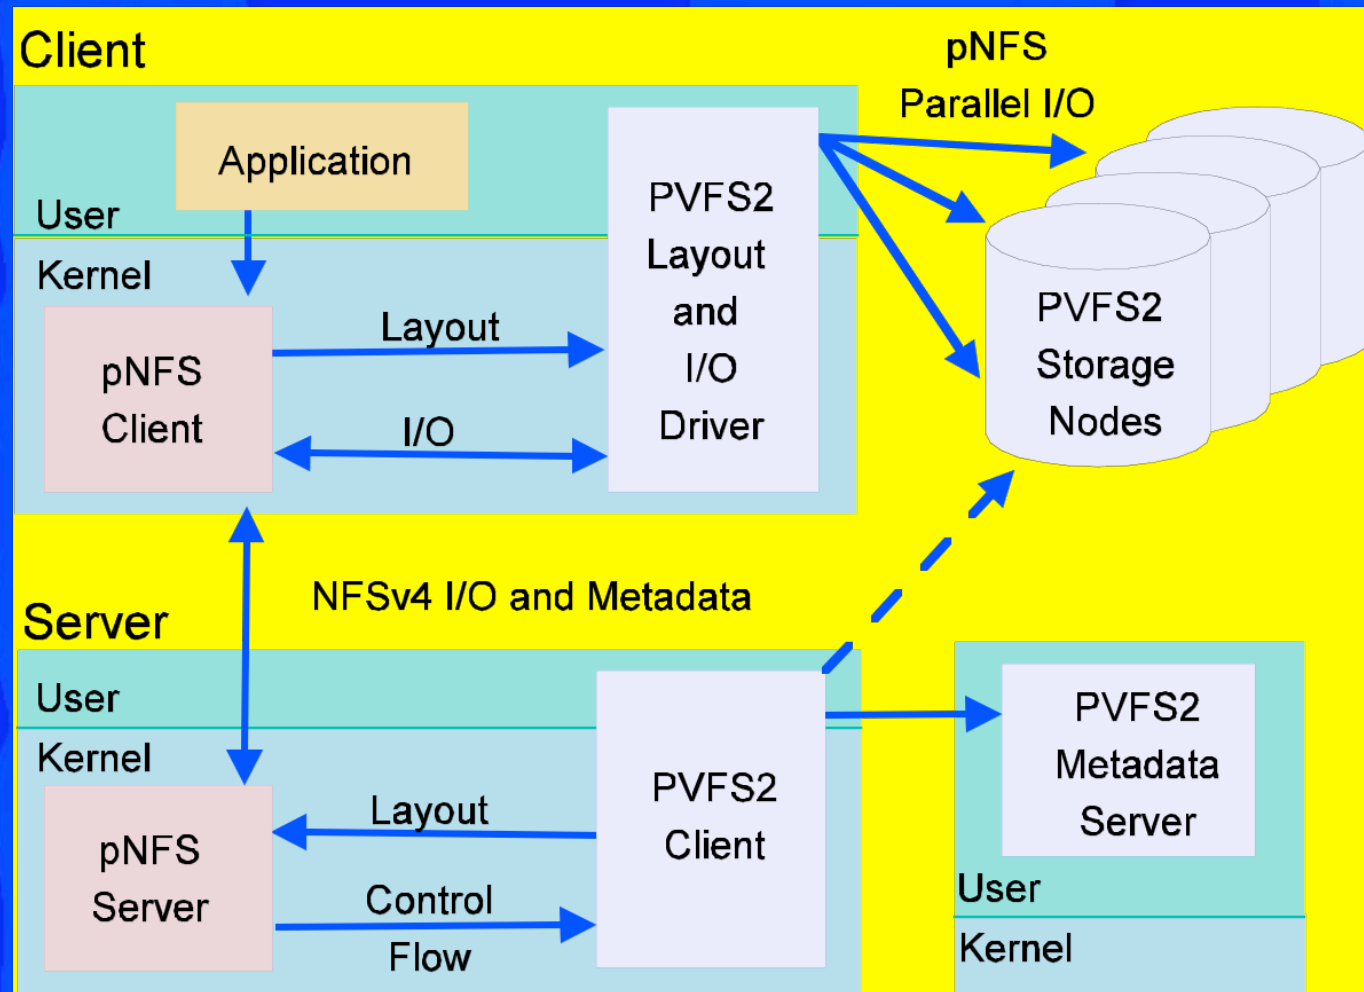


pNFS schematic

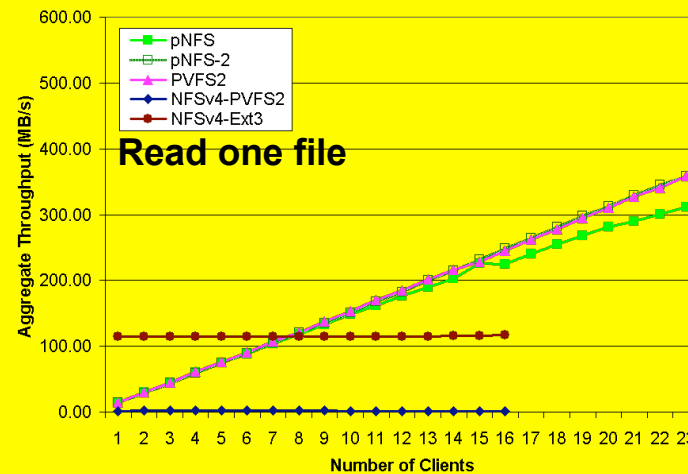

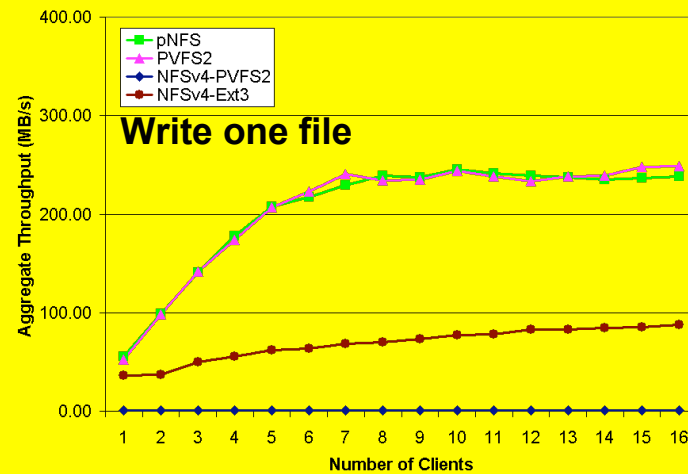
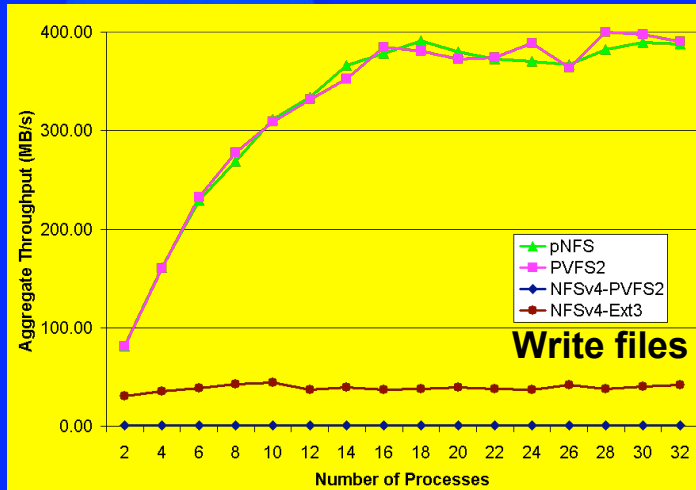
- ◆ Client accesses storage directly
- ◆ Bypasses NFS server to scale with storage



pNFS prototype on PVFS2



pNFS I/O performance



Access method comparison

