



October 12-14, 2004

NFSv4 Delegation: Past, Present, & Future

Bill Baker

Senior Staff Engineer

Sun Microsystems

Bill.Baker@Sun.com



October 12-14, 2004

NFSv4 Delegation

- Past: Protocol Overview
- Present: Solaris Implementation
- Future: Future Development



October 12-14, 2004

NFSv4 Delegation

Protocol Overview

What is a delegation?



A guarantee by the server that no conflicting access to a file will occur without notification.



October 12-14, 2004

NFSv4 Delegation

Protocol Overview: Goals for Delegation

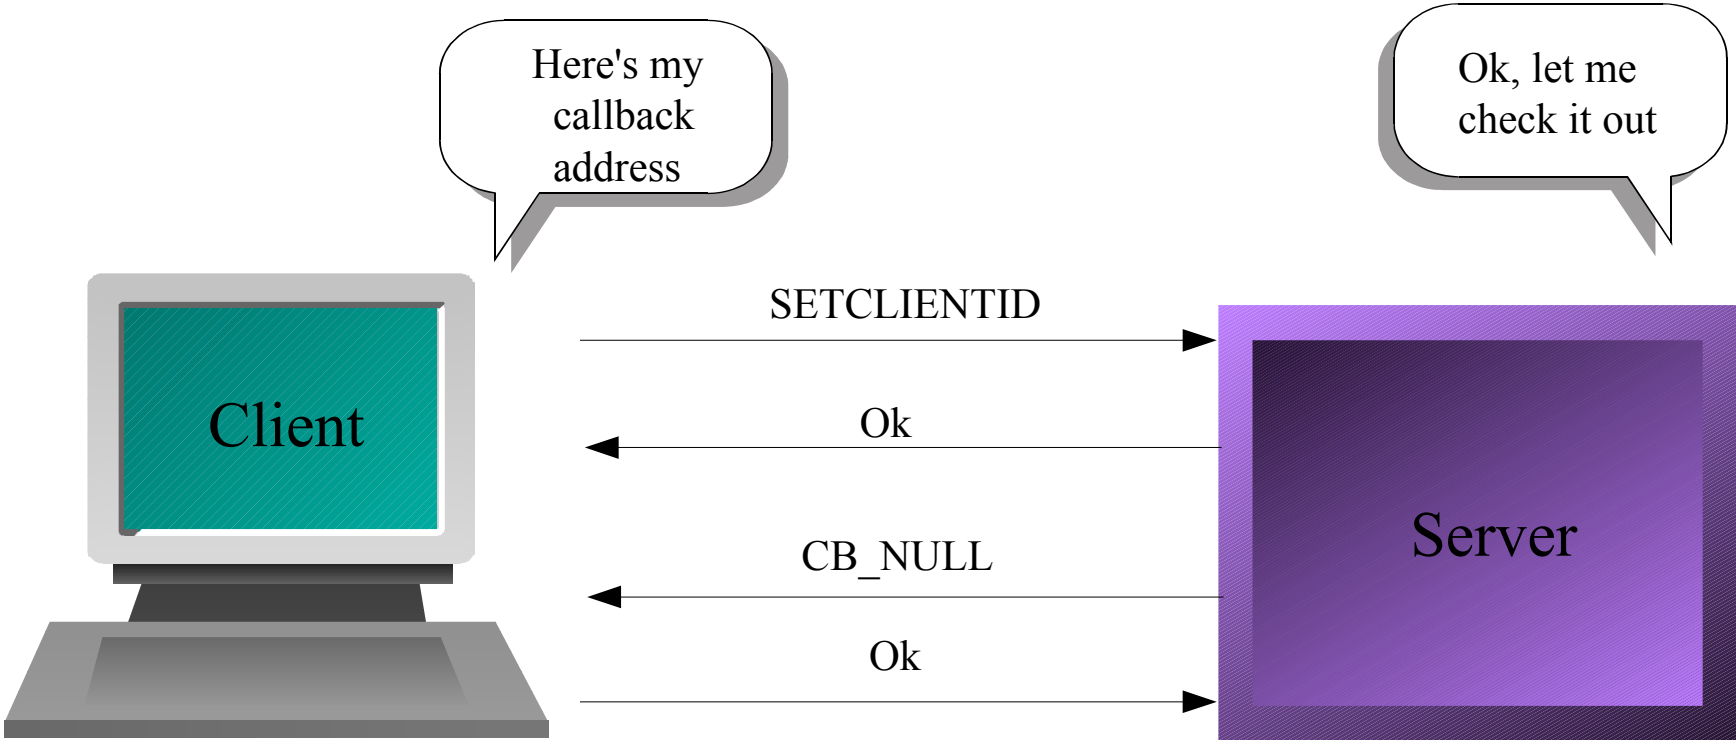
- Reduce number of “over the wire” operations
- Sweet Spots
 - Many clients reading
 - Single client writing
 - Repeated open/close operations



October 12-14, 2004

NFSv4 Delegation

Protocol Overview

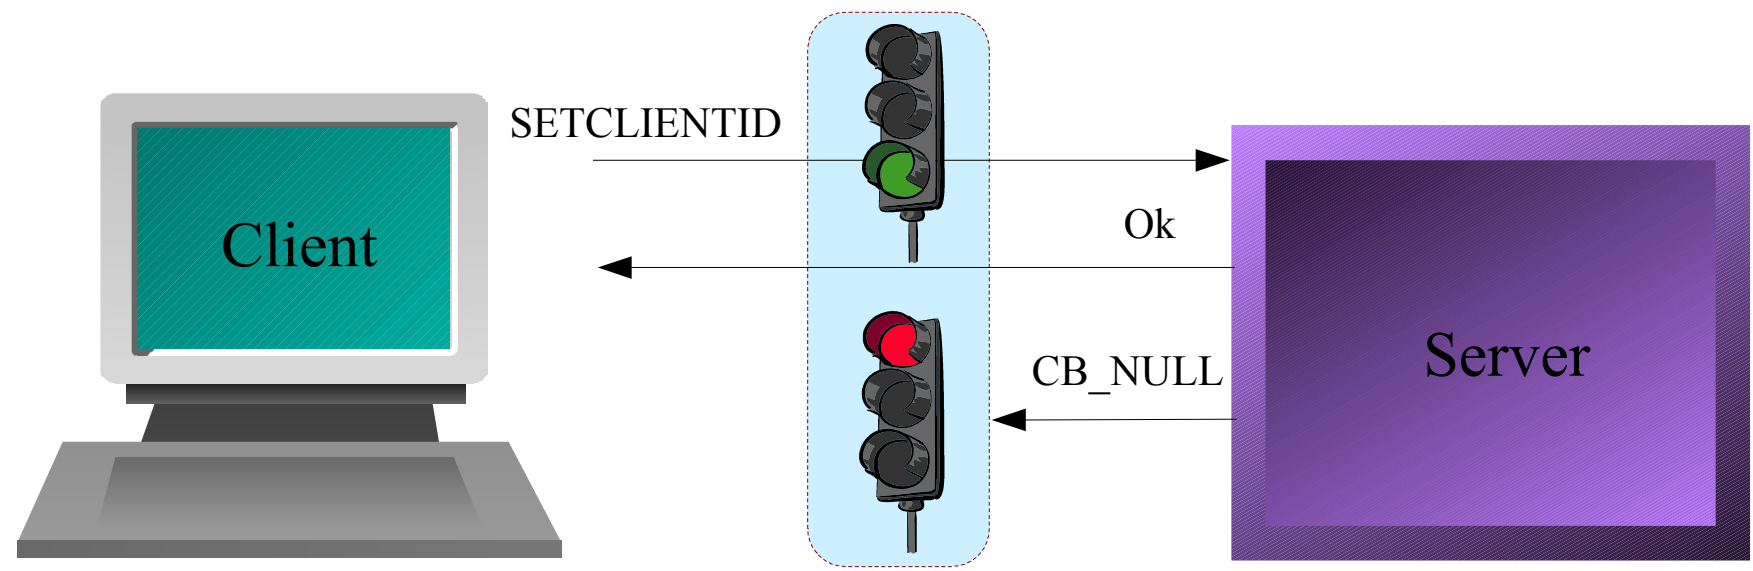




October 12-14, 2004

NFSv4 Delegation

Protocol Overview






October 12-14, 2004

NFSv4 Delegation

Protocol Overview




Server grants based on policy, heuristics & current access



October 12-14, 2004

NFSv4 Delegation

Protocol Overview



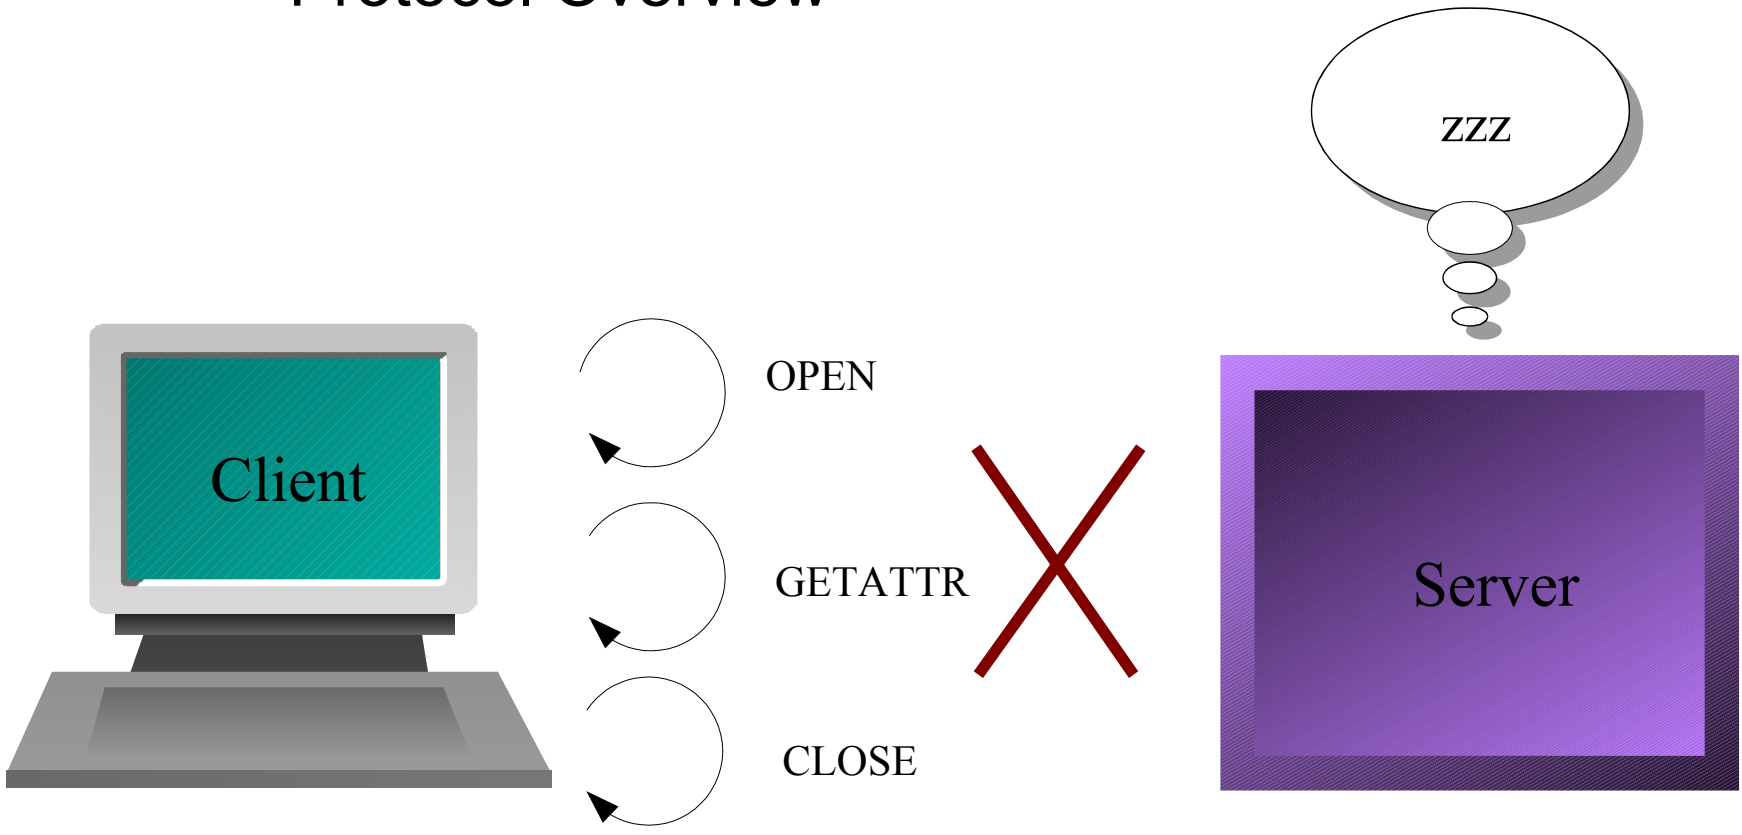
client retains the delegation after close of



October 12-14, 2004

NFSv4 Delegation

Protocol Overview



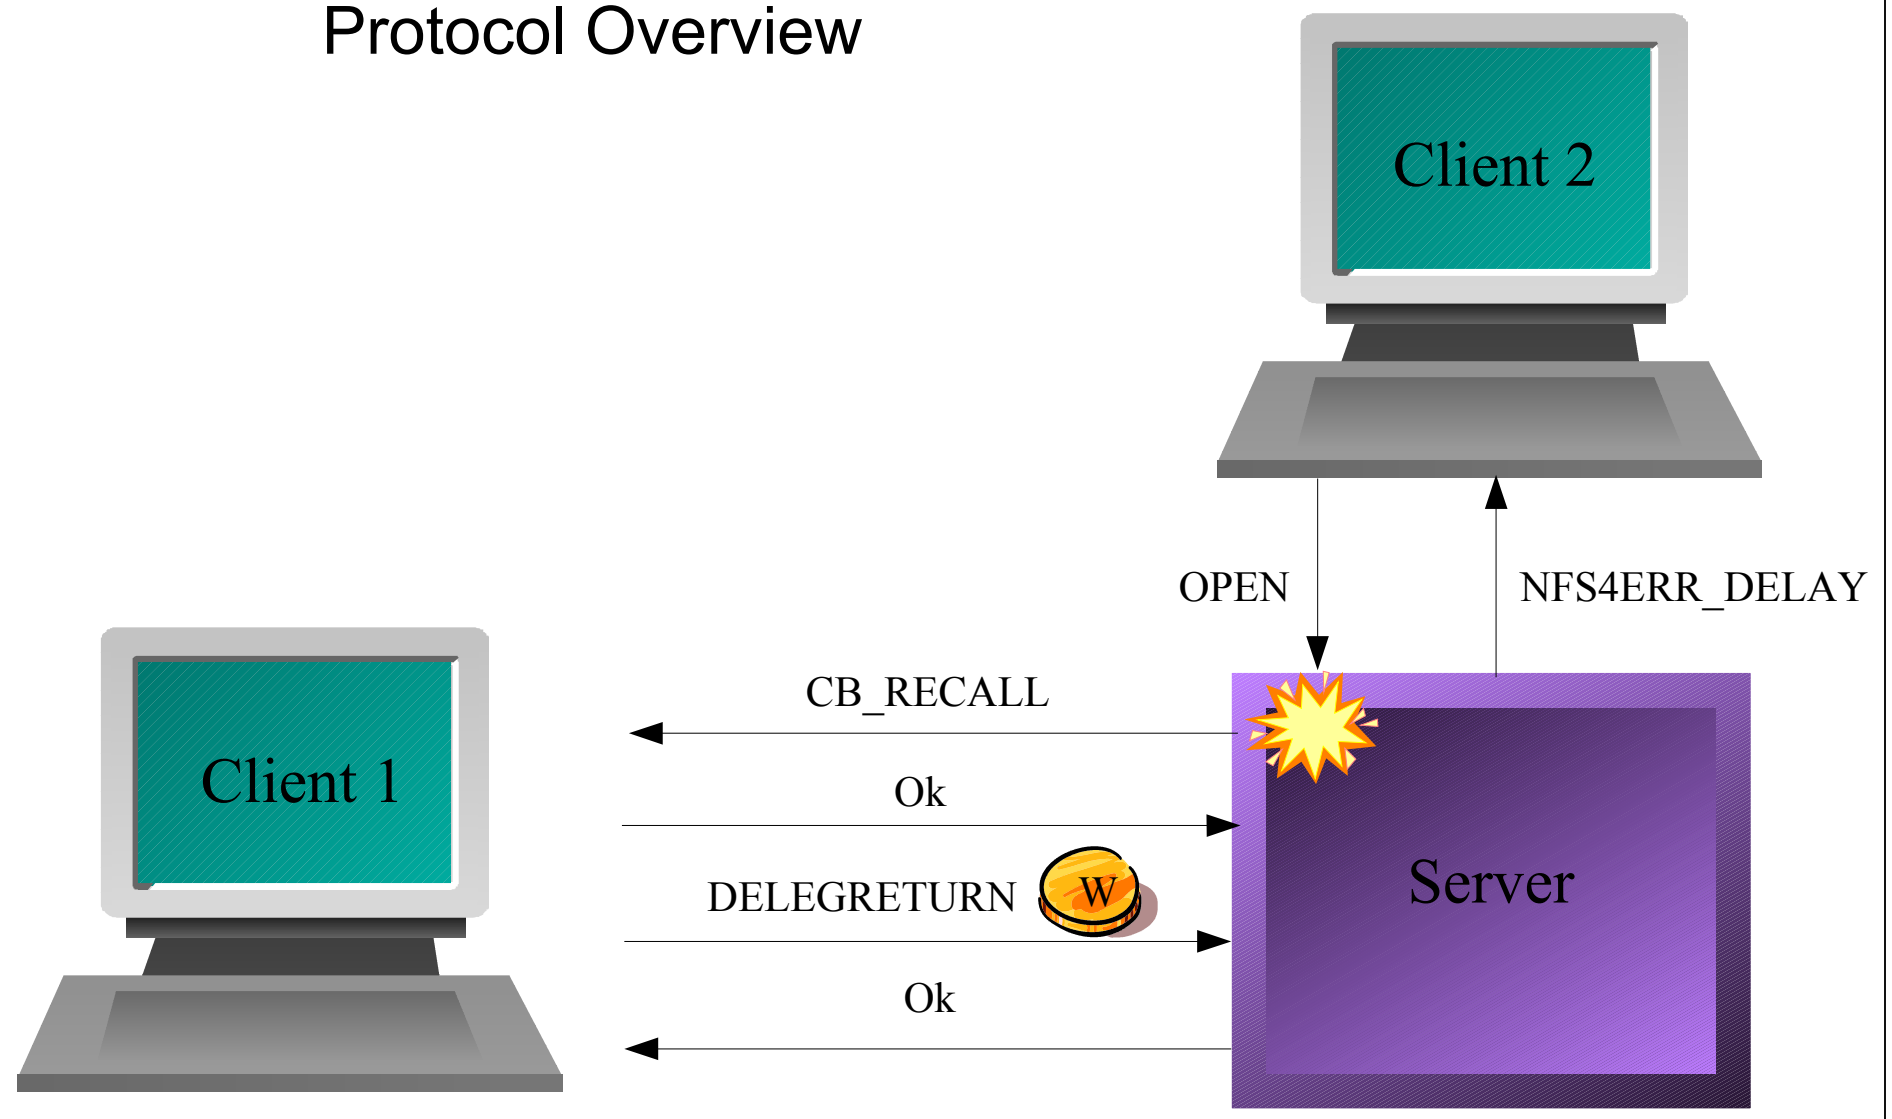
Client is freed from doing repeated getattrs



October 12-14, 2004

NFSv4 Delegation

Protocol Overview

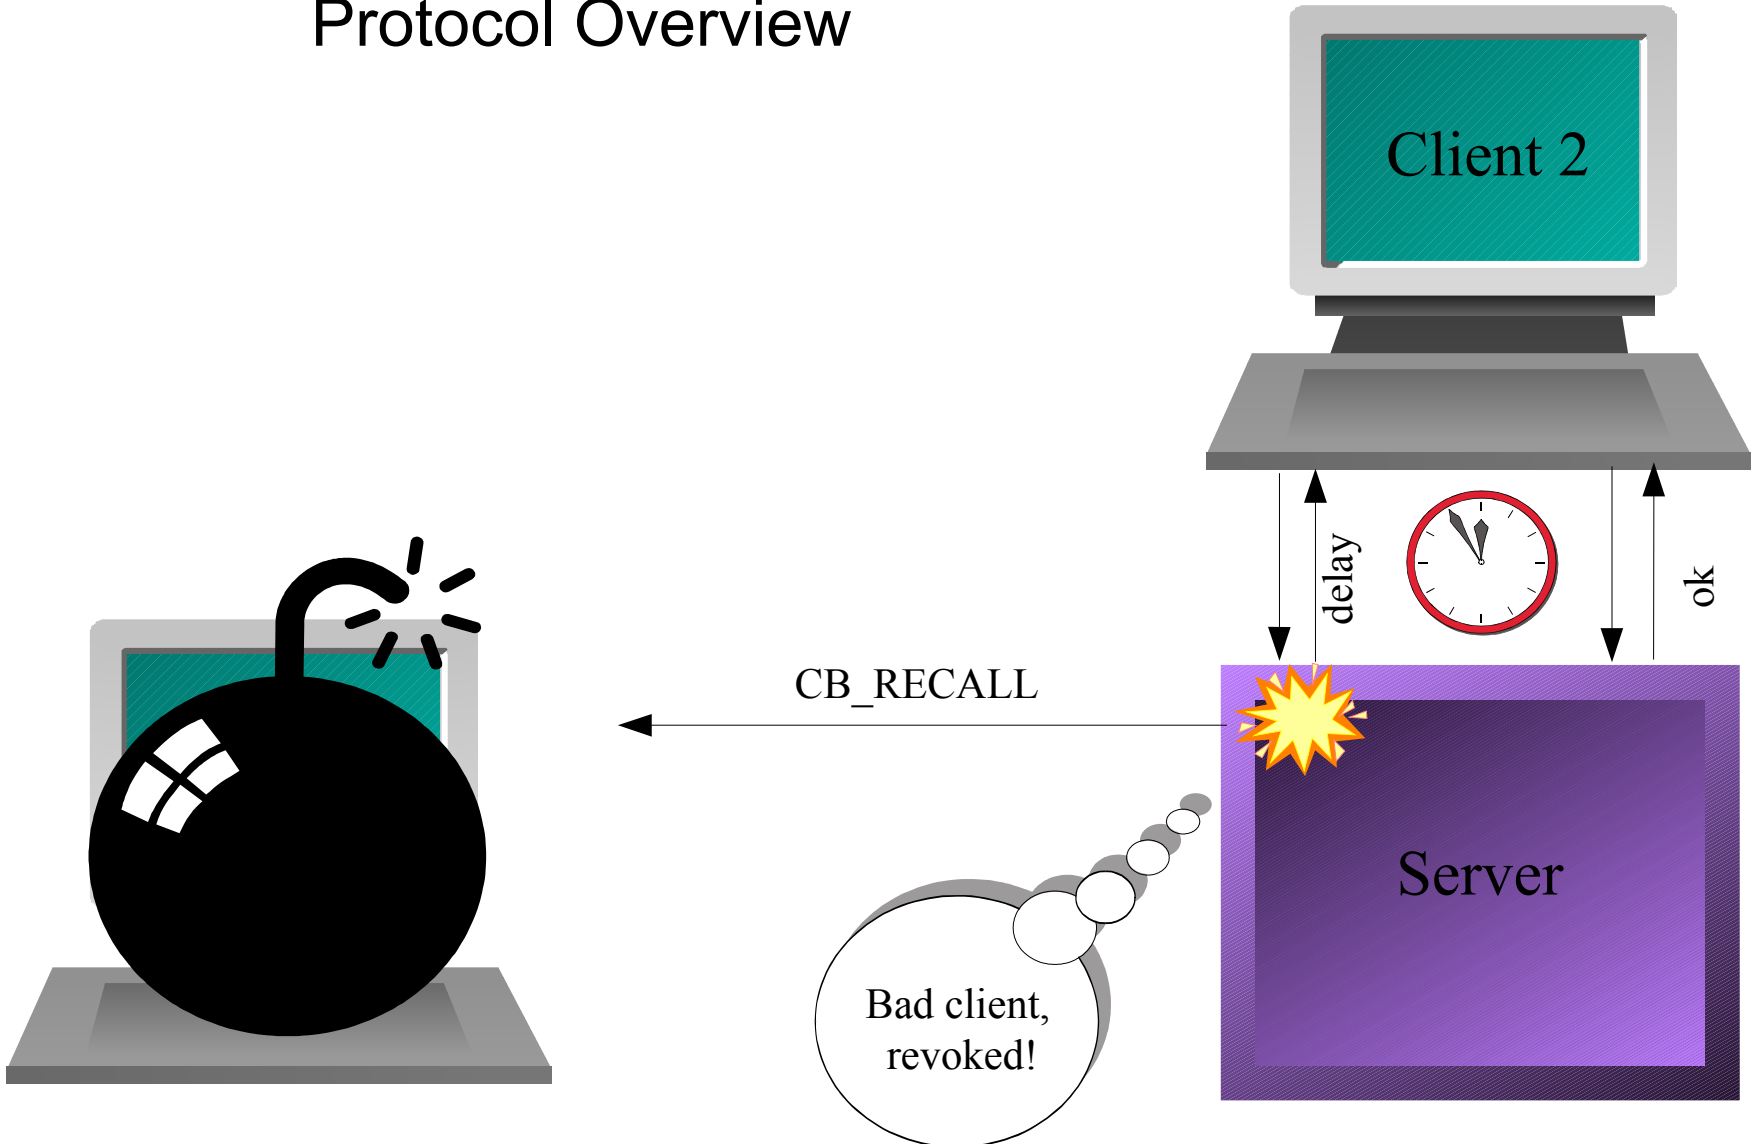




NFSv4 Delegation

October 12-14, 2004

Protocol Overview





October 12-14, 2004

NFSv4 Delegation

Solaris Implementation

- Enabled by default
- Server offers both read & write
- Integrated with local access
- Integrated with v2 & v3 access
- Disabled via `/etc/default/nfs`

```
NFS_SERVER_DELEGATION=off
```



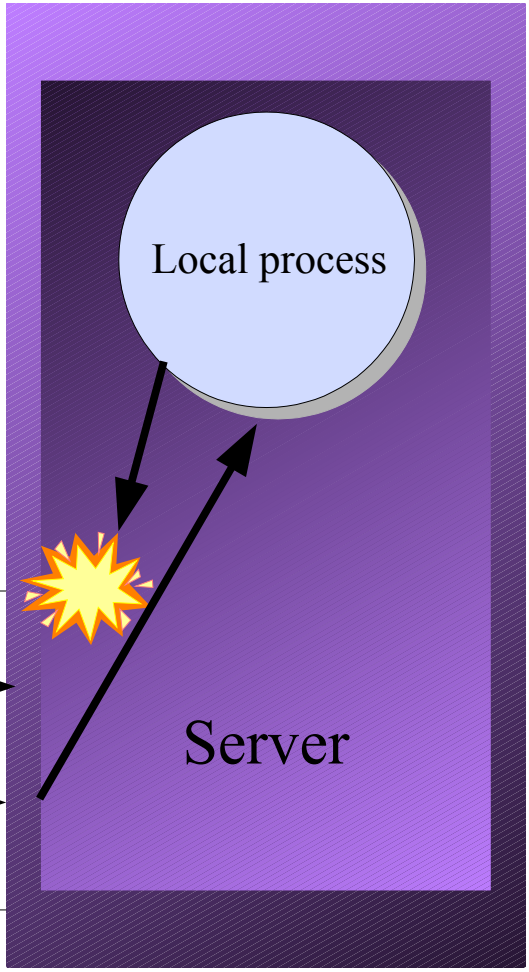

NFSv4 Delegation

October 12-14, 2004

Solaris Implementation

Local process opens file with conflicting access modes.

Local process is blocked until client returns the delegation

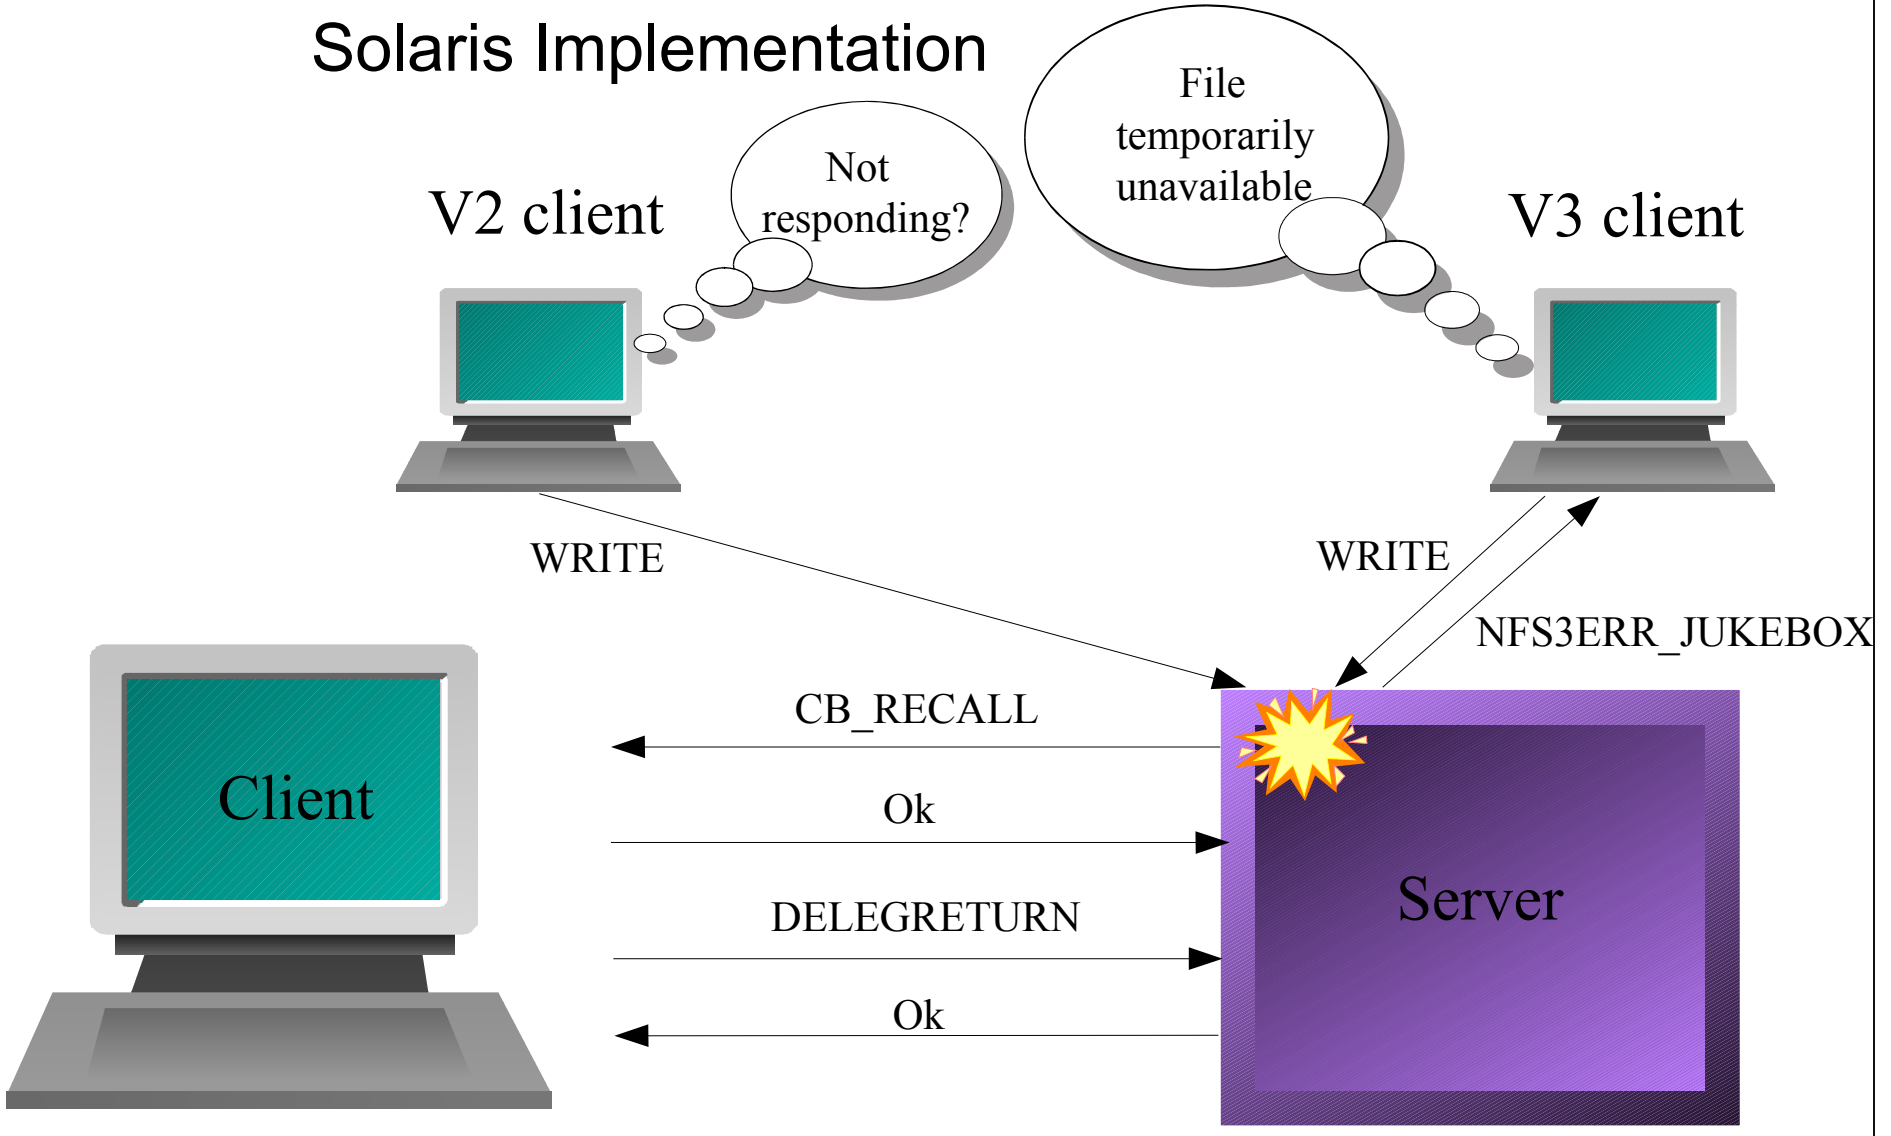




October 12-14, 2004

NFSv4 Delegation

Solaris Implementation





October 12-14, 2004

NFSv4 Delegation

Solaris Implementation, Performance Results

- **SPEC SFS not applicable**
 - Synthetic Client, protocol aware
 - Clients use compounds in different ways
- **Traditional Client Side Benchmarks**
 - Throughput tests – NO IMPACT
 - General tests

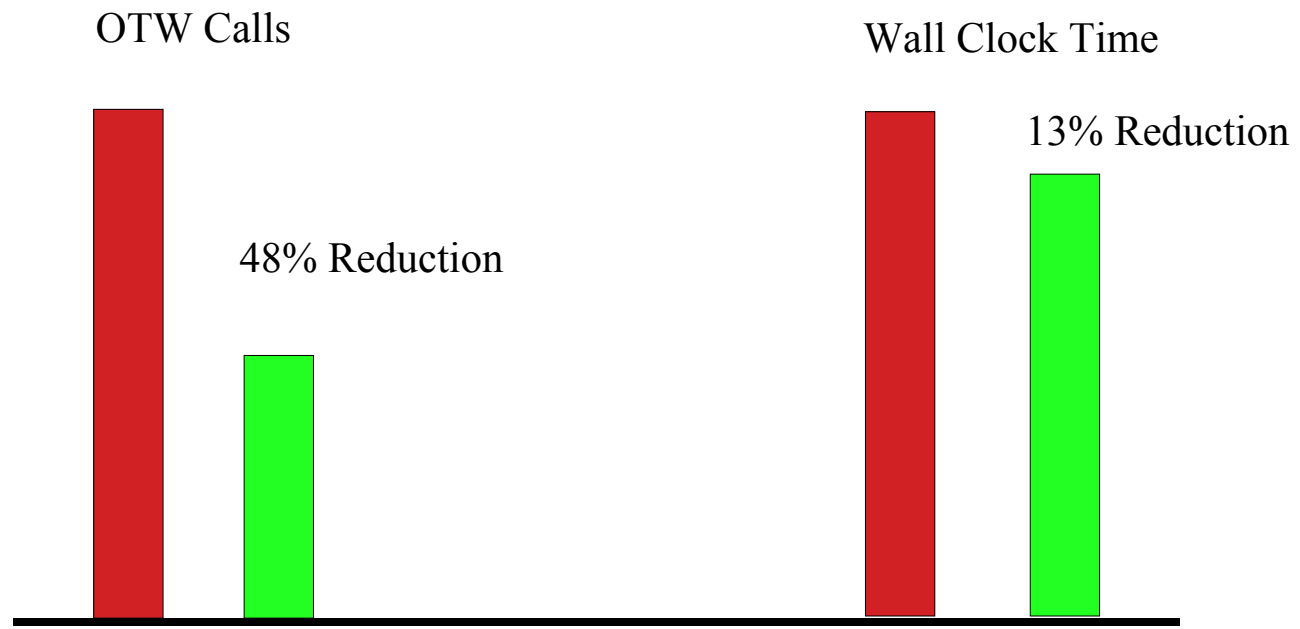


October 12-14, 2004

NFSv4 Delegation

Solaris Implementation, Performance Results

Benchmark: Full build of the Solaris OE over NFSv4

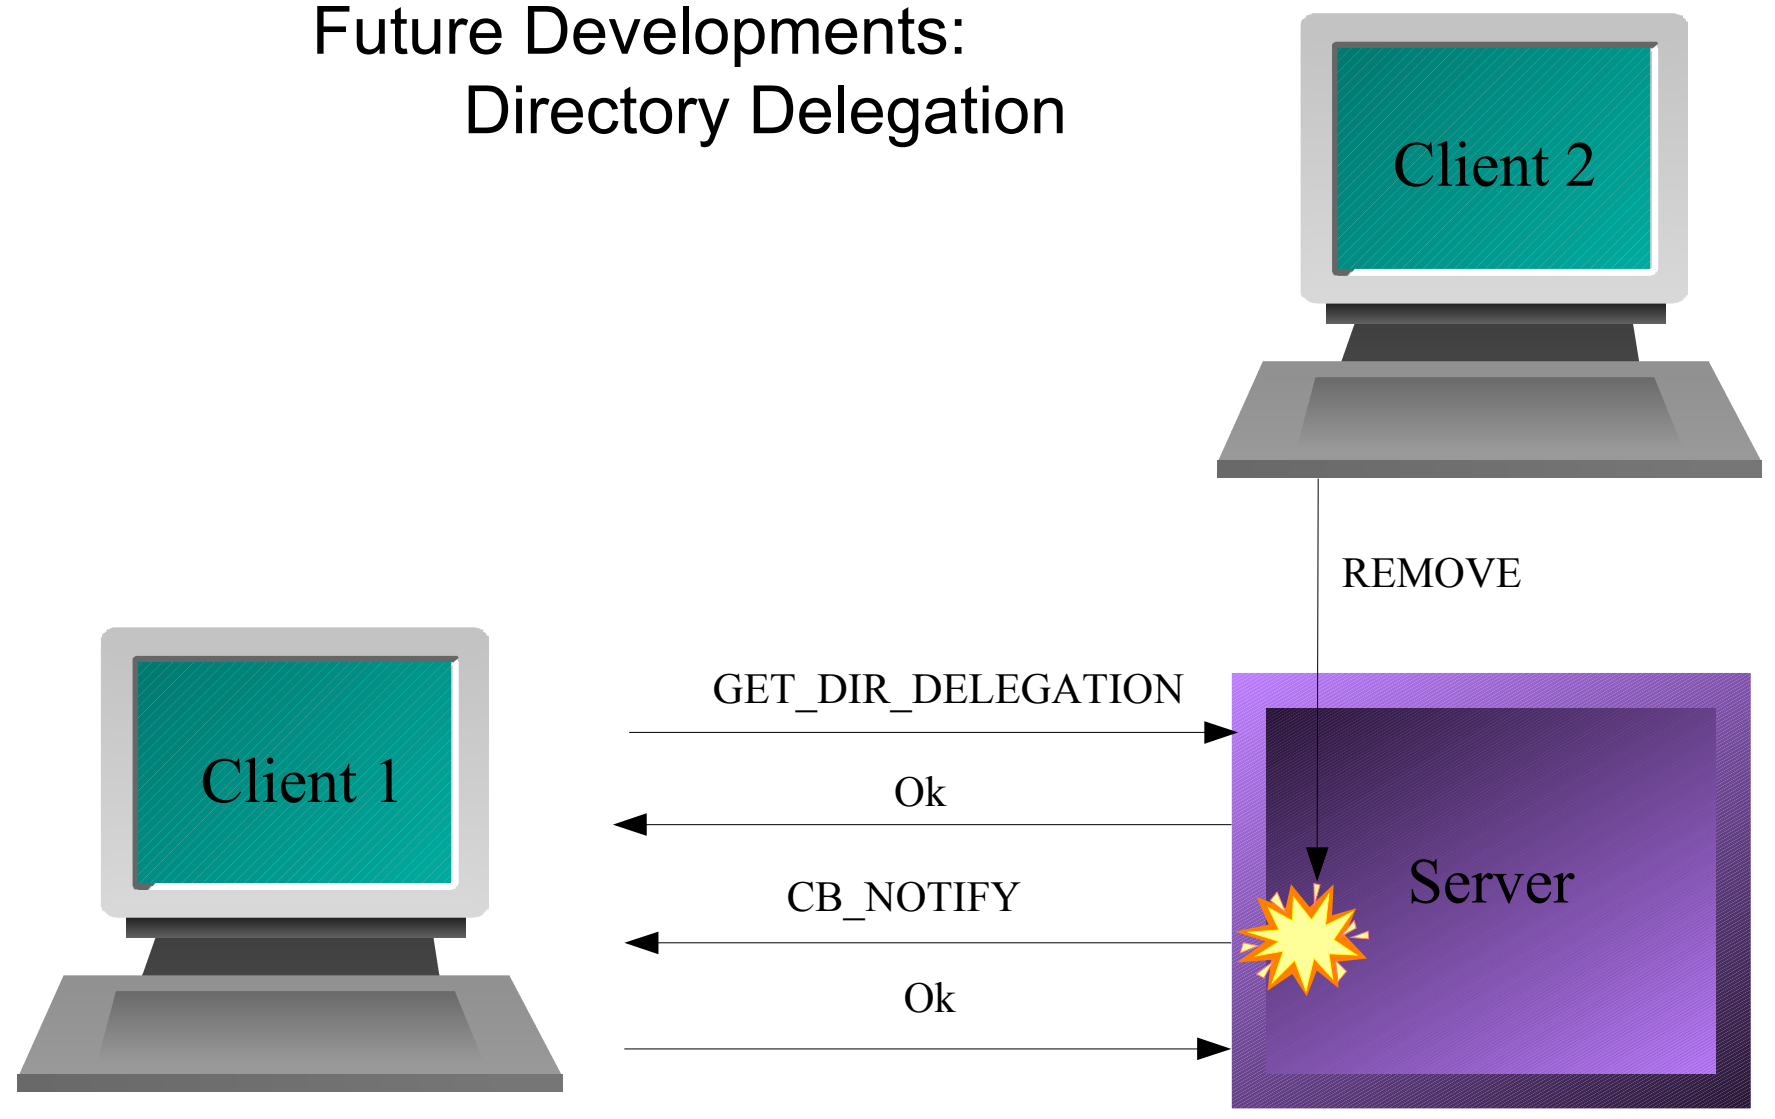




October 12-14, 2004

NFSv4 Delegation

Future Developments: Directory Delegation






October 12-14, 2004

NFSv4 Delegation

Future Developments:
Session Extensions



The callback channel can share the same connection as the operations channel.