

NFS Lock Manager

It finally works (I think)

Introduction

- NFS server is stateless (for v2 and v3)
- Record locks add state
- NLM protocol adds that state

How it works

- RPC requests for lock, test, unlock
- Blocked lock requests get a granted message when they get the lock
- Send a cancel message to interrupt a blocked lock request
- `rpc.statd` detects server and client reboots

Old rpc.lockd

- Implemented in userland
- Doesn't share state with local record locks
- Incomplete and in some parts broken
- Blocking locks can't be interrupted
- Semantics of flock(2) not supported

New rpc.lockd

- Implemented entirely in the kernel
- Supports everything except DOS shares
- New kernel-mode RPC infrastructure (both client and server)
- Initially implements NLM server, NLM client coming soon

Other stuff

- Local locking augmented to support async lock requests and remote locks
- Robust graph-based deadlock detection
- Enforced fairness for contended locks (no thundering herds)
- Regression tests

Future work

- NLM client with interruptible locks, correct semantics for flock(2)
- Client crash recovery
- Cache coherence
- Improve rpc.statd - see <http://www.connectathon.org/talks06/talpey-cthon06-nsm.pdf>

Thanks to:

- Isilon Systems
- Alfred Perlstein

Questions?