

Migration in NFSV4.0

What's been going on since the last cthon

David Noveck
February 21, 2012

A Capsule history

- We (mostly me) wrote this up for RFC3530
- Some people tried to implement it
- They had some problems
 - Told people in a talk at the last cthon
 - I missed that
- Discussed issues informally at IETF 81
- Decided to address in an individual draft
 - Group was suffering from BDSAN
 - Bis Document Schedule Anxiety Neurosis (not in DSM)
- Made some progress – will report on where we are

What are the issues?

- All concern transparent state migration
 - Failure to free migrated state on client reboot
 - Server reboots result in a confused lease situation
 - Multiple leases with the same `nfs_client_id4`. Ouch!
 - Clientid proliferation
 - Wart or neoplastic disease?
- How serious are these?
 - Depends on how heavily you want to use migration

Individual Draft

- Entitled *NFSv4.0 migration: Implementation experience and spec issues to resolve*
 - Authored by me, Piyush Shivam, Chuck Lever, Bill Baker
 - Makes proposals for spec corrections to address the issue
 - Some of the proposals are inherently controversial
 - Principally those dealing with client-string models and trunking
 - -02 published in mid-January
 - Been some prototype implementation work
 - We'll hear about that.

Non-uniform client string model

- Non-uniform client-string model is traditional
 - Client uses different identifying string for each server it talks to.
 - Don't try this at the airport :-)
 - RFC3530 makes it a "should," but not a "SHOULD".
- When state moves from server to server
 - Can't merge it because state belongs to two different clients
 - Server not told they are the same (and has no way to find out)

Going to uniform client string model

- When state moves from server to server
 - Can't merge it because (it seems) state belongs to two different clients
 - Server not told they are the same
 - So we have to hack or change models
 - Tried to allow both. Finally decided change had to happen for clients to support migration
- Traditionally, been told this is impossible
 - Reasons keep changing
 - Decided the issue had to be addressed