

New Migration Draft

How we got here, where we are now, and some thoughts on going forward.

David Noveck

NFSv4 Working Group Meeting at IETF84

July 30, 2012

Overview

- **Background:**
 - Multi-server features are harder than they look
 - Migration implementers and spec authors found that out
- **Work so far**
- **Latest Document**
- **Going Forward with migration fix-up**
 - In v4.0
 - In v4.1

Multi-server features: Needs

- One of the biggest changes in NFSv4, but ...
 - That wasn't really recognized
 - It didn't seem as big as *statefulness*, for example
 - Fundamental infrastructural needs were not promptly attended to:
 - Client instance definition (See Chuck's presentation)
 - Server identification (only explicitly addressed in NFSv4.1)

NFSv4 Multi-server features: List

- **Migration**
- **Replication**
- **Referral**
 - **FedFS**
- **pNFS**
- **Server-to-server copy**

Migration Implementation: Situation to be Dealt With

- **Major symptoms**
 - Failure to properly free state (on client reboot)
 - Clientid4 proliferation, leading to
 - Situations not foreseen by protocol
 - Client complexity
- **Situation we had:**
 - Feature as specified could be prototyped
 - But wasn't ready for real use

Migration Implementation: Underlying Issues

- No merging of state bundles for different fsid's that are between same pair of client/server instances
 - Spec doesn't tell you to do that
 - And if it had, you wouldn't know how to do it:
 - Cannot tell if two state bundles belong to *same* client, since that information is hidden from servers
- That problem goes back to recommendation (in RFC3530) to use non-uniform client strings
 - At least it wasn't a RECOMMENDATION 😊

Previous Documents

- **Co-authored with those doing implementations**
 - Chuck Lever, Piyush Shivam, Bill Baker
- **Started with individual drafts**
 - Three iterations of draft-dnoveck-nfsv4-migration-issues
 - From August 2011 to January 2012
 - Biggest change was to move toward RECOMMENDING use of uniform client-string model for clients supporting migration.
- **First working group draft in April 2012**
 - Pretty close to draft-dnoveck-nfsv4-migration-issues-02
 - Added some discussion on NFSv4.1

Latest Document

- **Draft-ietf-nfsv4-migration-issues-01**
 - Out three weeks ago
 - Please read it and comment
 - Preferably in that order :-)
- **Changes from the -00**
 - More attention to NFSv4.1
 - Address NFS4ERR_CLID_INUSE
 - Better address issues raised by older servers for clients using uniform client strings

Going Forward on Migration: Overall

- **Expect at least one more draft of existing informational document**
- **At some point, we will have to decide about effecting needed changes in standards-track documents**
 - **Issues are different for NFSv4.0 and NFSv4.1**
 - **See next few slides**

Going Forward on Migration: For NFSv4.0

- **Scope of changes:**
 - Well-defined with lots of proposed text 😊
 - Pretty extensive 😞
 - Won't fit in errata; Updates are more bis-like
 - But the 3530bis train is all assembled 😊
 - But bad for migration stuff 😞
- **Need to address this *somehow*:**
 - A 3530tris is conceivable, but won't get off the ground
 - Expect a migration-specific bis-like working-group update document
 - Other possibilities?

Going Forward on Migration: For NFSv4.1

- **Scope of changes:**
 - Smaller than v4.0 – No `nfs_client_id4` issue
 - But still seems too big for errata
- **Other differences:**
 - 5661bis train just forming up 😊
 - Prototyping work for 4.1 migration likely to be delayed 😞
- **Would like to put into 5661bis.**
 - Have to balance:
 - Bis schedule
 - Prototyping expectations
 - Quality/completeness/review expectations