

Referrals in NFSv4

Dave Noveck
Connectathon '05

Six W's and an H

- Referrals – What?
- Referrals – Whence?
- Referrals – Why?
- Referrals – How to
- Referrals – Where?
- Referrals – When?
- Referrals – Whither?

What are Referrals?

- A way of directing a client to a filesystem
 - Usually on a different server
- How does it differ from migration, since,
 - Migration directs a client to another fs
 - Usually on a different server
- So does a referral, but ...
 - It happens when a client, accesses the fs after it moved

Referrals vs. Pure Referrals

- Referral is a per-client concept
 - Client B can see a referral, after
 - Client A sees a migration ... Or,
- All clients can see a referral
 - Migration happens before any client access
- All clients *must* see a referral
 - Pure referral case
 - Fs is long gone or was never really there

Where do referrals come from?

- From RFC3530 of course
 - But RFC3530 does not mention referrals
 - True, but you have to look between the lines
- How closely between the lines?
 - Do you need an electron microscope? *No.*
 - Do you need a light microscope? *No.*
 - Do you need a magnifying glass?
 - Yes and a careful look at the protocol logic

From Migration to Referrals

- Referrals are a limiting case of migration
- Start with general migration case
 - Keep moving migration event back in time
 - There is a set of fh's migrated (volatile ones go away)
 - Earlier migration and the set gets smaller
- Two interesting cases
 - Singleton set: Only fh you have is root
 - Null set: No fh's in fs and you have a referral

Transition to Referrals

- Suppose we have:
 - PUTROOTFH/LOOKUP/GETFH
 - LOOKUP crosses into new fs
 - Migration event happens somewhere in there
- If it happens just after GETFH,
 - No error now. MOVED when you use it.
- 1 us earlier (just before GETFH)
 - MOVED on GETFH and you have a referral
- 1 ms earlier (before request)
 - Same thing – referral
- Years earlier (maybe before RFC3530 was written :-)
 - Can't use v4 fh's for this fs – pure referral case

Why implement referrals

- Above leads to question:
 - That's kind of cool/weird but why do that?
- If you implement migration,
 - Need limiting case
 - Clients not synchronized with migration event
- Pure referral case:
 - Allows multi-server (even global?) namespace
 - May evolve to AFS-style namespace in NFS clothing
 - What else would be needed to get there?

How to implement

- But RFC3530 doesn't explain this so how can one implement it?
- Glad you asked:
 - Read draft-noveck-nfsv4-referrals-00.txt
 - Implementation Guide to Referrals in NFSv4
 - Dave Noveck, Carl Burnett
 - Implement that, test at next Bakeathon
- Seriously, draft seems to work for us
 - But we need comments from others, especially comments based on implementation experience
 - We'll try to keep document up-to-date

Where are they now?

- In ONTAP, server prototype (a bit hacky)
- In AIX, client prototype
- In AIX, server prototype
- Linux implementations
 - Client and server prototypes
 - Being adjusted for interoperability

When???

- Now for prototypes
- For products, not sure
- First successful intra-vendor referral
 - AIX prototype in January, 2005
- First successful inter-vendor referral
 - AIX client and Netapp server on March 1, 2005
 - AIX client and Netapp/IBM servers on March 2, 2005
- Early days
 - But day 1/2 feels a lot better than day –x
 - Not there yet but we are somewhere

Where do we go now?

- Improve implementations
 - More testing
 - Improve code (i.e. hack removal)
- More implementations
 - Who's interested? You will have something to test against. Next Bakeathon maybe.
- Improve document
 - Comments please
 - Working group draft? Informational RFC?
- Start thinking about a central namespace defn