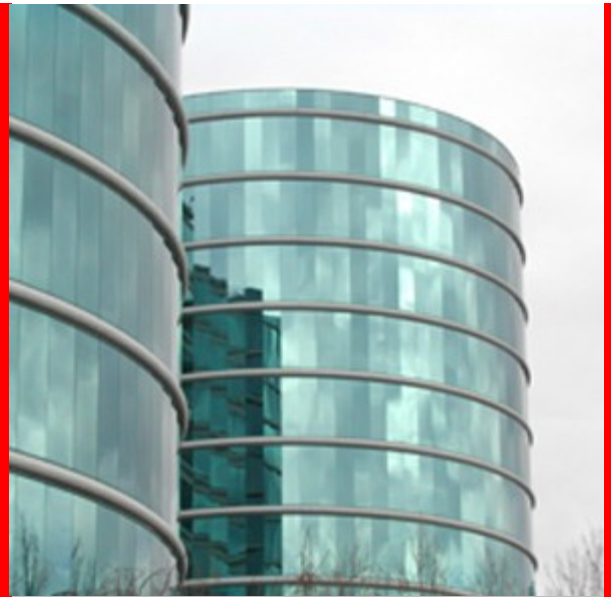

Sun

ORACLE

ORACLE THIRTY YEARS 30

Federated FS 2010 update

- Rob Thurlow
- Solaris NFS Group

From last year – my brief history

- First, there was replication and failover
  - ... with lists of locations a client could follow
- Then came migration and a MOVED error
  - "It's not here anymore, see the list"
- Then we extended MOVED to referrals
  - "It's not here (and never was, actually)"
- Then we baked it into NFSv4.1
  - And sometimes even tested it
- Now, we need to manage server referrals
  - That's FedFS!

FedFS Motivation

- We want to build uniform namespaces
  - /nfs4/oracle.com/home/thurlow same across the whole enterprise
  - Based on managing NFSv4.x referrals
  - With any type of hardware or OS, of course!
  - Making use of existing server contents
  - With no changes to NFSv4 protocol
- How's your implementation coming?

FedFS Protocols

Namespace Management	Namespace Navigation
1) NSDB Management (LDAP)	3) Client root discovery (DNS)
2) Junction Management (ONC RPC)	4) Junction resolution (LDAP)

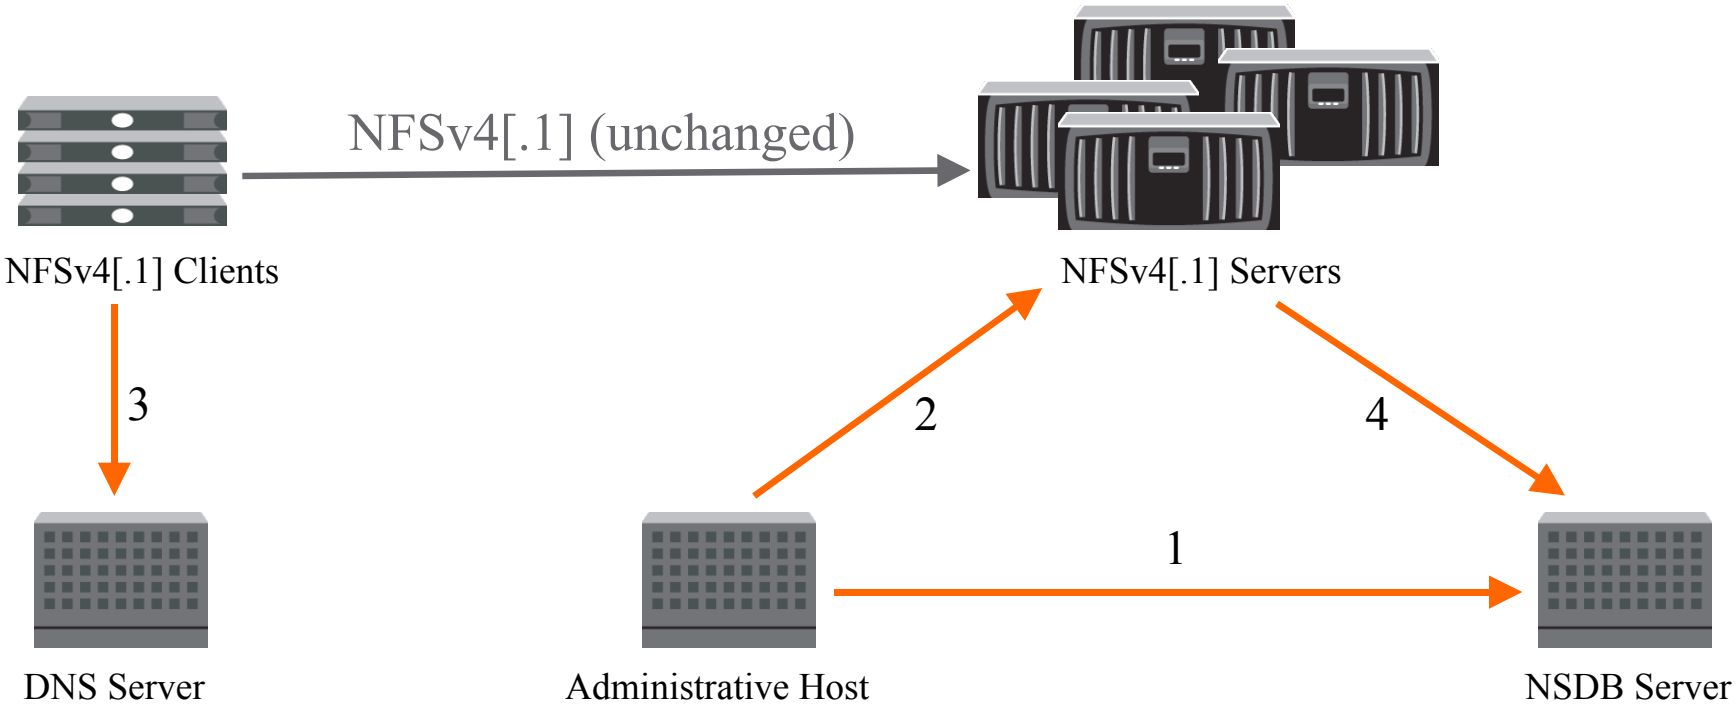

Documents In Play

- "Using DNS SRV to Specify a Global File Name Space with NFS version 4"
  - http://www.ietf.org/id/draft-ietf-nfsv4-federated-fs-dns-srv-namespace-02.txt
- "Administration Protocol for Federated Filesystems"

>http://www.ietf.org/id/draft-ietf-nfsv4-federated-fs-admin-04.txt

- "NSDB Protocol for Federated Filesystems"

>http://www.ietf.org/id/draft-ietf-nfsv4-federated-fs-protocol-05.txt

Since Last Year ...

- NFSv4 WG charter added FedFS 3/09!
- Two Bake-a-thons, three IETFs
- Maintained SNSDB sample code
- Requirements published as RFC 5716!
- Survived DNS directorate review and LDAP expert review
- Three protocol documents near last call!

Decisions Made

- NSDB can now be on a dedicated or shared LDAP server
  - > NSDB Container Entry points within DIT
- NSDB schema uses XDR for pathnames for better locale support
- Admin protocol has LOOKUP_FSN for diagnosis
- Admin protocol has ops to manage “trust anchors” for NSDB security

Decisions Made

- Conversion tools (e.g. from automounter maps) don't need to be standardized
- File servers may cache NSDB to TTL
- Clarified details of how client finds a namespace

Root Fileset

- We plan to specify this optional feature
- For now, we envision this:
  - A special root fileset can be provided
  - Based on special LDAP entries
  - Instantiated when a root fileset server boots
  - Read-only and very simple
    - No ACLs, files or named attrs

Futures

- Last call coming soon
- Building and testing prototypes
- Root fileset work next
- Replication and migration work coming

Another Admonition!

- Get involved!
- Read RFC 5716 to orient yourself
- Read the drafts and make comments
- Join the weekly call
  - Invites always on NFSv4 WG alias
- Build and test implementations

Appendix

Glossary

- Fileset: a directory tree that can be managed
- FSN (fileset name): a unique fileset identifier
  - UUID plus an NSDB name (see below)
- FSL (fileset location): network location of a fileset instance
  - Server name plus relative path
- Junction: an object on a server that stores an FSN
- NSDB (namespace database): LDAP service that tracks mapping between FSNs and FSLs