



NFSv4 ID Mapping

Rick Mesta
Connectathon 2005
Feb 28 2005



Topics

- NFS Identifiers; then, now, ...
- Multi-domain challenges
- Case study: Sun's WAN
- Solaris' approach: DNS TXT RR
- NFS4 Domain selection in Solaris
- IETF Proposal(s)

NFS Identifiers: then, now...

- Past:
 - v2/v3 == 32 bit otw attrs
- Present:
 - RFC 3530: `[user | group]@domain`
- Future:
 - General purpose string/id mapping service

Multi-domain challenges

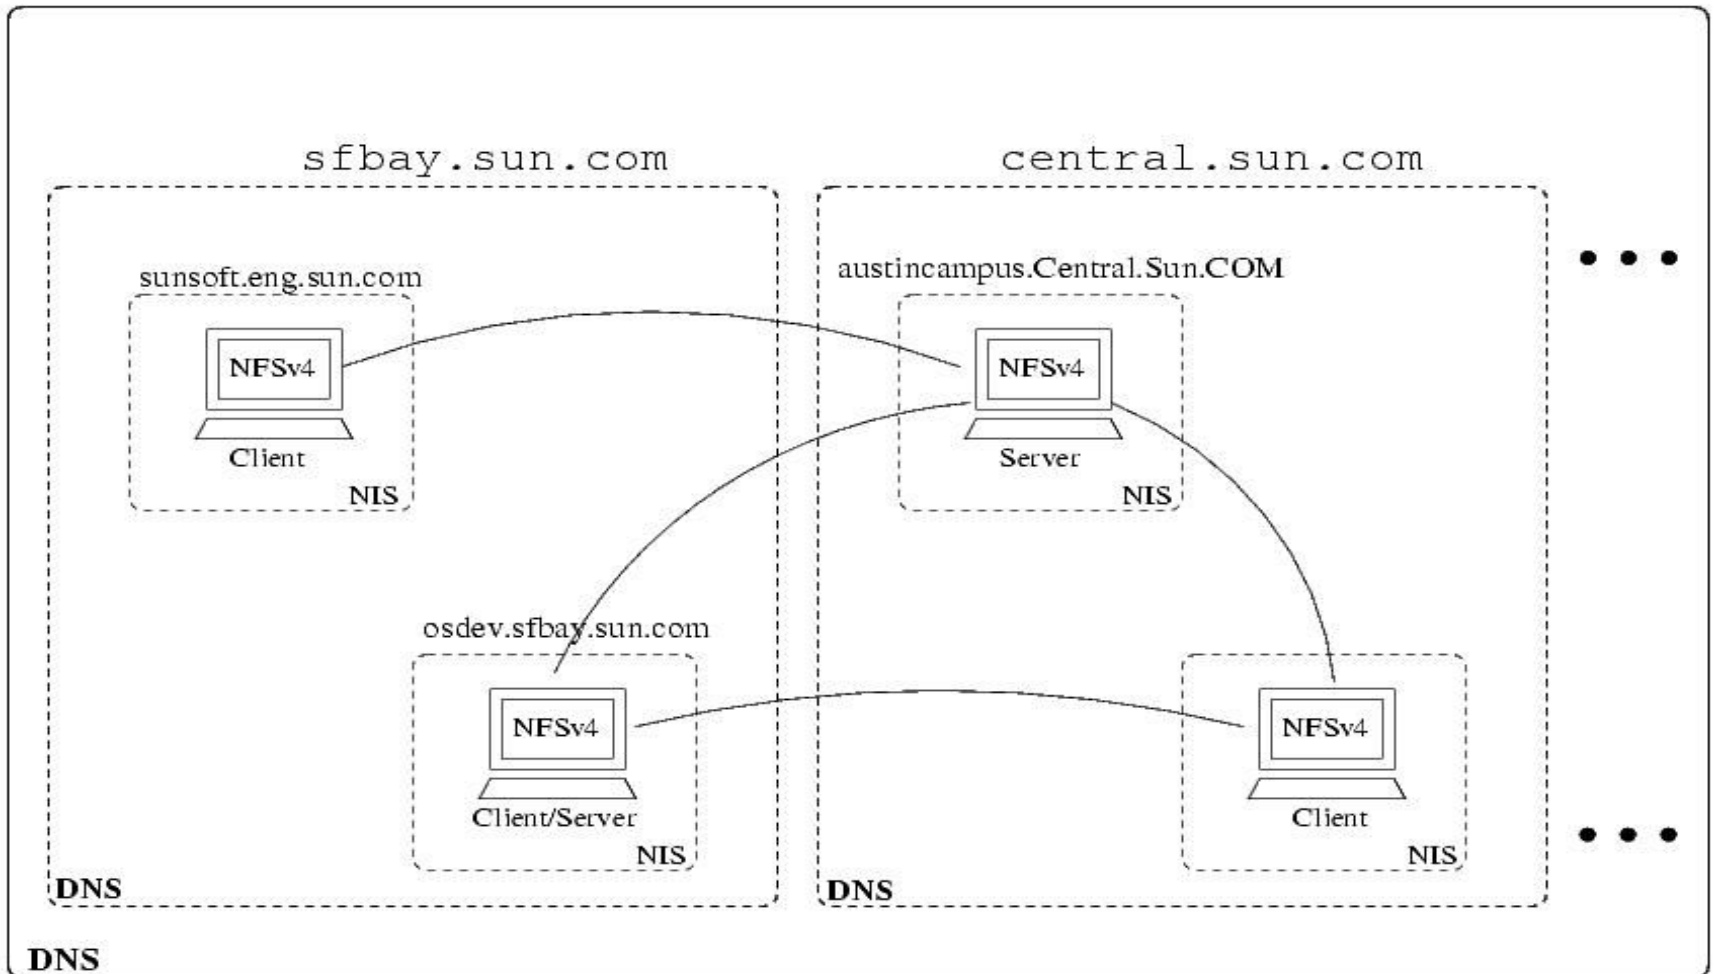
- The Good:
 - file/dir ownership from RPC/krb info (CREATE)
- The Bad:
 - Unix OS's don't support domain qualified id's
 - SETATTR failures on mismatched domains
 - “nobody” reporting storms on client
 - ripple-effect problems on apps that `stat(2)` files to determine ownership (ie. gnome)

Case Study: Sun's WAN

- One corporate level domain “sun.com”
- Constitutes multiple-geo's:
 - Typically, one sub-DNS domain per geo
 - Geo sub-DNS domain has 1[+] NIS domain(s)
 - Lab systems typically configured NIS-only

Sun WAN: Typical NFS4 traffic

sun.com



Case Study: Sun's WAN cont'd

- Key Points

- Typical NFSv4 access is across domain boundaries
- Client/Server domain mismatches cause:
 - “nobody” reporting storms (regardless of id space)
 - `owner/owner_group SETATTR` failures
- Solaris does not natively handle domain qualified id's

Solaris impl: DNS TXT RR

- Pros
 - DNS is ubiquitous
 - High probability of use for hostname resolution
 - Can leverage as NFS4 domain distribution svc
- Cons
 - Additional sysadmin intervention
 - TXT RR's are problematic due to sub-typing
 - Not the final/gen purpose id/string mapping svc

Domain Selection in Solaris

- `/etc/resolv.conf`
 - Does DNS TXT RR exist ?
 - `TRY_AGAIN`: Fire requery thread; Move on...
 - `NETDB_SUCCESS`: Yes; Done !
 - `NO_DATA`: No! Move on...
 - Use DNS domain itself
- `domainname (1m)`
- `/etc/default/nfs (NFSSMAPID_DOMAIN)`

IETF Proposal(s)

<http://www.ietf.org/internet-drafts/draft-mesta-nfsv4-domain-01.txt>

- **Motivation**

- Unix OS's don't handle domain qualified id's (yet)
- Multi-domain NFSv4 deployment is problematic
- Mismatched domains cause *nobody* storms
- DNS likely to be used for hostname resolution

- **Interim-Solution**

- TXT RR's are susceptible to sub-typing
- DNS NFS4 specific RR I-D to be submitted
- Stop gap measure until gen purpose string/id mapping svc can be adopted

Summary

- Multi-domain NFS4 envs are problematic
- No common way to configure domain (as per RFC 3530)
- DNS is likely to be used
 - leverage it
 - can mitigate problem until general string/id mapping svc available
 - TXT RR's not an elegant solution; NFS4 specific RR I-D to be submitted



NFSv4 ID Mapping

rick.mesta@sun.com

