

ORACLE®



## **FedFS Domain Root: Observations**

Chuck Lever  
Principal Member of Staff



# Overview

- Feature review
- Read-write domain root
- Domain root availability
- SMB support

# FedFS Domain Root

## Salient Features

- Top-level directory of cross-server namespace
- Contains NFSv4 referrals pointing to actual content
- How to find a domain root on the network:
  - DNS SRV records advertise servers exporting domain root
  - Well-known export path
    - Read-only replica: */domainroot-example.net*
    - Read-write replica: */domainroot-write-example.net*
- How to find a domain root on an NFS client:
  - Globally-useful names:
    - Read-only replica: */nfs4/example.net*
    - Read-write replica: */nfs4/.example.net*

# FedFS Domain Root

## Read-write domain root

- NFSv4 clients do not manage NFSv4 referrals
  - When would NFS clients need to mount a read-write replica?
- Client pathnames similar to AFS, but not quite the same
  - AFS: writability controlled by volume name
  - NFS: writability controlled by client mount options and inheritance
    - Clients must mount read-only domain root replica read-write so that users still have full access to submounts
    - Clients might today use “ro” to enable more aggressive caching (use case: shared “/usr”)



# FedFS Domain Root

## Domain root availability

- Built-in availability features
  - DNS SRV RRs have weight and priority attributes
    - Client sorts these when it mounts domain root
    - SRV information generally not available in kernel context
- Thus clients with already-mounted domain roots are vulnerable to reboots of servers exporting domain root

# FedFS Domain Root

## Future SMB support

- No FedFS protocol indication to clients whether to mount SMB or NFS version of FedFS domain root
- SMB “//share/path” names do not look like NFS globally-useful names: “/nfs4/domain/path”
- FEDFS\_LOOKUP\_JUNCTION cannot today return an FSL that is not a fedfsNfsFsl



**ORACLE IS THE INFORMATION COMPANY**