

ORACLE®



## **Linux Client-Side Migration Status**

Chuck Lever <[chuck.lever@oracle.com](mailto:chuck.lever@oracle.com)>

Principal Member of Staff

3<sup>rd</sup> party logo

## Migration Outline

- Minor version 0 behavior
- Minor version 1 behavior
- Delivery schedule
- Ongoing testing



## Migration Recovery Architecture

- Migration recovery done in existing state recovery logic
- Separate recovery thread started as needed
- Typically NFS task sleeps while recovery runs
- Callouts handle recovery specific to minor version

# NFSv4.0 Migration



## Handling NFS4ERR\_MOVED

- If op is not LOOKUP, assumes migration
- NFS task sleeps while recovery thread runs
- Source: [ PUTFH; GETATTR(fs\_locations); RENEW ]
- Client attempts to connect a transport to each server in the returned fs\_locations list
- Client clones long-form and short-form client ID
- Destination: [ RENEW; SETCLIENTID ]
- Performs state recovery if needed

## Handling NFS4ERR\_LEASE\_MOVED

- NFS task sleeps while recovery thread runs
- Loops over FSIDs associated with this client ID
  - Sends [ PUTFH; GETATTR(type) ]
  - If server returns NFS4ERR\_MOVED, try migration recovery
  - Repeat until server stops returning NFS4ERR\_LEASE\_MOVED

# NFSv4.1 Migration





## Handling NFS4ERR\_MOVED

- If op is not LOOKUP, assume migration
- NFS task sleeps while recovery thread runs
- Source: [ SEQ; PUTFH; GETATTR(fs\_locations) ]
- Client attempts to connect a transport to each server in the returned fs\_locations list
- Destination: [ BIND\_CONN\_TO\_SESSION ]
- Performs state recovery if needed

## Handling SEQ4\_STATUS\_LEASE\_MOVED

- NFS task **does not** sleep while recovery thread runs
- Loops over FSIDs associated with this client ID
  - Sends [ SEQ; PUTFH; GETATTR(type) ]
  - If server returns NFS4ERR\_MOVED, try migration recovery
  - Repeat until server clears SEQ4\_STATUS\_LEASE\_MOVED

# Delivery Schedule



## Rough Integration Plan

- Code review and integration - June/July 2011
- Merge with 3.1 - August/September 2011
- 3.1 becomes final - Thanksgiving 2011?
- UEK lags upstream 12-18 months
- RHEL 6 (and OEL 6) plans unknown

# Migration Test Plan



## Ongoing Testing

- Testing with a non-TSM server
- Multiple client ID recovery
- Testing with GSS and security negotiation
- Upstream means wider availability
- NFSv4.1 is unproven until we have a server

## Migration Recovery, Missing Pieces

- Reconnecting NFSv4.1 clients should try EXCHANGE\_ID
- We're not careful yet about using the correct credential for RENEW and SETCLIENTID
- fs\_locations\_info not yet implemented
  - No sorting of returned fs\_locations
  - No support for read, write, readdir, file handle classes
- Volatile file handles not supported

# Q&A

ORACLE®



# Hardware and Software

ORACLE®

## Engineered to Work Together

ORACLE®

ORACLE®