

# **Current Updates to NFSv4 Specs**

## ***Trunking and Other Material Related to Multi-Server Namespace Features***

**David Noveck (davenoveck@gmail.com)**

**Nfsv4 WG meeting at IETF105 in Montreal**

**July 25, 2019**

# Summary

- NFSv4.0 Changes
  - Adds trunking discovery to NFSv4.0.
  - Published as RFC 8537 (Proposed Standard)
- NFSv4.1 changes
  - Adds trunking discovery, Transparent State migration to NFSv4.1
  - Not published as RFC yet (Long story, starting at Slide 6)
- Implementation discussion

# NFSv4.0

## Protocol Description Updates and Status

- Trunking for NFSv4.0
  - Adds trunking discovery to the trunking detection defined by RFC 7931.
- Defines interaction of trunking and migration
  - Support for changes in trunking patterns with and without migration.
- Document status
  - Published as a Proposed Standard
  - Happened by the Milestone date <sup>^^</sup> , although one can argue about “final” submission

# NFSv4.1 Protocol Description Updates

## Major Elements

- Adds description of trunking discovery, building on the trunking detection mechanism in RFC 5661.
- Discusses interaction of trunking and migration.
  - Includes support for changes in trunking patterns with and without migration
- Adds description of transparent state migration, paralleling the transparent state migration in NFSv4.0, but adding
  - Session migration
  - Migration of pNFS metadata server

# NFSv4.1 Protocol Description Updates

## Necessary Cleanup

- Revised use of the term “replica”
  - Needed to distinguish clustered server treatment from trunking.
- Clarified handling of per-fs grace periods.
  - Provides continuity of lock state w/o transparent state migration.
- Revised descriptions needed for two ops:
  - Existing EXCHANGE\_ID inconsistent with state migration.
  - RECLAIM\_COMPLETE wasn't clear about handling of per-fs case
- Cleanup of error descriptions to accommodate per-fs grace periods and interaction of trunking and migration.

# NFSv4.1 -- Document Progress and Status

## Finding out what IESG Really Didn't Like

- Original Document followed, “here are the updates to apply to the spec” paradigm.
  - OK during WGLC, IETF last call, AD review, but ....
  - The IESG would not accept it, feeling that the update model is too hard on readers.
- Clearly, we needed a more bis-like document
  - But a full bis is a big undertaking (at least a year away).
  - Settled on draft-ietf-nfsv4-rfc5661-msns-update.
    - Had a complete revised multi-server namespace section.
    - Also new EXCHANGE\_ID, RECLAIM\_COMPLETE, and error changes.
  - IESG still didn't like it.

# NFSv4.1 -- Document Progress and Status

## Trying to Find out what IESG Would Like.

- IESG wanted an even more bis-like document.
- But a full bis would take a lot of time
  - Issues with RFC8178, RFC8434, Internationalization, Security
  - Details are in later talk.
- Finally settled on a partial bis
  - Document obsoleting RFC5661
  - Has replacement section for multi-server namespace plus associated changes to enable those.
  - But no other changes to RFC5661.

# NFSv4.1 -- Document Progress and Status

## The Moment of Truth

- Produced draft-ietf-nfsv4-rfc5661sequi-msns.
  - -00 submitted on 6/25/19
  - Needs to go through WGLC etc. before it is considered by IESG
- Need people to read and comment
- Will submit -01 next week
- At that point, we should decide whether WGLC is appropriate.
- After that, would like to move document forward

# Implementation Issues

## Items for Continuing Discussion

- Trunking Implementation
  - NFSv4.0 is published so work could begin, unless one is v4.0-averse.
  - Work on NFSv4.1 would, formally speaking, be a prototype until RFC publication, even though substantive change now is not likely.
  - Given similarity in trunking for the two versions, it makes sense to work on them together.
  - Discussion, probably informal, is necessary to co-ordinate client and server work.
- Transparent state migration implementation
  - Implementations of NFSv4.0 state migration could be upgraded.
  - Apparently, some prototype work has already been done.