

# **Multi-server Namespace in NFSv4.x**

**Previous and Pending Updates**

**Dave Noveck**

**WG Meeting at IETF 102**

**July 16, 2018**

# Overview

- Multi-server namespace was an important advance for NFSv4.x, but...
  - It wasn't addressed very well in RFCs 3530, 5661, or 7530 †
  - As author/editor of those msn chapters, all I can say is “Mea Culpa”.
- Work has been going on to address those issues:
  - RFC7931 (Proposed Standard as of 7/2016)
  - draft-ietf-nfsv4-migration-issues (will discuss; hope to get to WGLC soon)
  - draft-ietf-nfsv4-mv0-trunking-update (at -00; Proceeding OK; Will discuss)
  - draft-ietf-nfsv4-mv1-msns-update (at -01; Proceeding OK; Will discuss)

# Msns Issues

## Areas to Address

Area	Description	Work Now Needed
Trunking Detection	Given two IP addresses, a means of determining whether they are connected to the same server.	None; OK as is
Trunking Discovery	Means of determining a set of alternate IP addresses that might be used to access the same server as an existing IP address.	Needed for v4.0 and v4.1
Transparent State Migration	Means by which two cooperating servers can effect a migration of state in a fashion transparent to the clients, eliminating the need to reclaim locks.	Needed for v4.1.
Multiple Connection Types	Describes the handling of multiple connection types (e.g. TCP and RDMA) reachable at the same address are dealt with and the proper connection type decided upon.	Needed for v4.0; Clarification needed for v4.1.

# Msns Issues

## RFCs and Standards Track Documents

Document	Trunking Detection	Trunking Discovery	Transparent State Migration	Multiple Conn. Types	Status
RFC3530	No. Trunking is a problem to be avoided.		Mentions.	Not mentioned.	PS. Obsoleted by RFC7530.
RFC5661	Yes.	Missing.	Missing.	Discussion incomplete.	PS. Needs to be updated.
RFC7530	No. Trunking is still a problem to be avoided.		Discussed but treatment inadequate.	Not mentioned	PS. Updated by RFC7931.
RFC7931	Provides as an option for v4.0.	Still missing	Addresses issues for v4.0.	Not Mentioned.	Proposed Standard.
mv0-trunking-update	Leaves as is in RFC7931.	<b>Provides for v4.0.</b>	Leaves as is in RFC7931.	<b>Will address.</b>	Submission target: 11/2018.
mv1-msns-update	Leaves as is in RFC5661.	<b>Provides for v4.1.</b>	<b>Provides for v4.1.</b>	<b>Will clarify and complete.</b>	Submission target: 3/2019.

# Draft-ietf-nfsv4-migration-issues

## Summary

- Informational document to discuss and guide needed work.
  - Could publish as RFC, but it is not clear if we want to
  - Working group needs to decide (See Process-beyond-WGLC Slide)
- Addresses a lot of stuff:
  - Migration, including transparent state migration.
  - Trunking detection and discovery
  - Multiple connection types
- Deals with issues in v4.0 and v4.1
  - Most problems are the same but often the solutions are different.

# Draft-ietf-nfsv4-migration-issues

## State of the document

- Material in -14
  - Trunking discovery (for both protocols)
  - Transparent state migration (retrospective for v4.0; prospective for v4.1)
  - In pretty good shape
    - No Pending issues
    - Needed changes reflected in draft-ietf-nfsv4-{mv0-trunking, mv1-msns}-update-00.
- Changes in -15
  - Handling of multiple connection types
  - Discussion of possible extensions (to v4.2)
- Original target date has gone by; WGLC should be by 8/2018.

# Draft-ietf-nfsv4-migration-issues

## Getting to WGLC

- Complete working group review and discussion.
  - For material in -14, review should already have occurred.
  - If it hasn't, it needs to be done soon.
  - Changes in -15 have been out for a month; might need further review.
- Get to a version on which WGLC can be requested.
  - Expect it to be -16, which will include material about the rca\_one\_fs issue discussed in [Next Slide](#).
  - Please provide comments on -15 by 8/10, three weeks from end of IETF102.
  - Based on comments, will produce -16 as a basis for WGLC.

# Draft-ietf-nfsv4-migration-issues

## rca\_one\_fs Angst

- As Rick Macklem has discovered:
  - Vmware client routinely does RECLAIM\_COMPLETE with rca\_one\_fs set to TRUE although FALSE is in line with the intention of RFC5661.
  - RFC5661's discussion of RECLAIM\_COMPLETE is not explicit as it might be.
  - Servers accept the current Vmware client behavior (ignoring the flag issue)
- Such server behavior is not particularly troublesome when fs migration does not occur.
- But it does impact the ability of servers to implement fs migration, while supporting such clients.
- Draft-ietf-migration-issues will need to discuss the potential conflict.

# Draft-ietf-nfsv4-migration-issues

## Process beyond WGLC

- After end of WGLC, document could be submitted for publication.
  - But it isn't clear if it is worth doing that.
  - Working group needs to decide.
- Some pro's:
  - Explains the overall reasons for change in way RFC7931 and draft-ietf-nfsv4-{mv0-trunking,mv1-msns}-update can't.
  - Also contains material about possible NFSv4.2 extensions.
- Some con's:
  - Haven't published such documents in the past.
  - Not clear who would read it once it is published

# Draft-ietf-nfsv4-mv0-trunking-update

## Summary

- Fulfills the v4.0 needs identified in migration-issues doc.
  - Similar to mv1-msns-update.
  - But much of the contents of that document is covered, for v4.0, in RFC7931
- Has any needed changes in this area, for v4.0.
  - Adds trunking discovery (very similar to mv1-msns-update).
  - Trunking detection already present, as an option in RFC 7931.
  - -01 expected to have further material on multiple connection types (analogous to work added in mv1-msns-update-01).

# Draft-ietf-nfsv4-mv0-trunking-update

## Further Work Needed

- Further working group discussion needed:
  - Of the core material in -00
  - And also of changes expected in -01 (for multiple connection types)
- Following needed for final document submission:
  - Submission of -01 (expected soon).
  - Going through WGLC.
- Target date of 11/2018 looks doable

# Draft-ietf-nfsv4-mv0-trunking-update

## Process Going Forward

- Discussion of existing prototype implementations would be helpful.
  - Given the schedule, discussion of these should happen soon.
- Review will take place on wg mailing list:
  - If people think a conference call would be helpful, one can be scheduled.
  - Could be combined with discussion of mv1-msns-update

# Draft-ietf-nfsv4-mv1-msns-update

## Summary

- Fulfills the v4.1 needs identified in migration-issues doc.
  - Parallels work done in RFC 7931 and mv0-trunking-update.
- Has any needed changes in this area, for v4.1.
  - Adds trunking discovery (very similar to mv0-trunking-update).
  - Trunking detection already present in RFC 5661
  - Adds transparent state migration to v4.1 (parallels RFC 7931).
  - -01 has further material on multiple connection types (analogous to work needed in mv0-trunking-update).

# Draft-ietf-nfsv4-mv1-msns-update

## Process Going Forward

- Further working group discussion needed:
  - Of the core material in -00
  - And also of the changes in -01 (out about one month)
- Prototype implementations would be helpful.
  - Given the 3/2019 target date, there is time for such work to contribute to document.
- Review will take place on wg mailing list:
  - Keep those cards and letters coming.
  - If people think a conference call would be helpful, one can be scheduled.

# Draft-ietf-nfsv4-mv1-msns-update

## Issues with regard to target date

- Current target is 3/2019 for final document submission
  - Implies need for WGLC by 1/2019
  - -01 is out, ready to be reviewed, and could be the WGLC candidate.
  - There is time for a -02, which might be necessary to address issues in rca\_one\_fs Angst slide.
  - Versions beyond that might require an adjustment of the date.