

Draft-ietf-nfsv4-rfc5661bis-01

Contents, Discussion, Future Steps

David Noveck (as document author)

For nfsv4 WG Meeting at IETF-117

July 28, 2023

Overview

- Motivations
- Contents
 - What's in it.
 - What is yet to be done.
- Topics for Discussion
 - Things done in -01
 - Things mentioned in -01, but not yet done.
 - Anything else people want to raise
- Process going forward

Motivations

What we started with

- Need to separate out areas that were common to all minor versions
- Tons of errata reports
- Need to Incorporate the material from documents updating rfc8881 (e.g. rfc8434)
- Unsatisfactory Treatment of Security in existing Spec
 - No Threat Analysis.
 - Discussion of AUTH_SYS presents it as OPTIONAL

Motivations

Related Documents (Slide One of Two)

- Internationalization -06
 - WG document, , now having ARTART review.
 - Reviewer assigned
 - Target date of 7/14 has not been met.
 - Need to figure out next steps.
- Security
 - Need to discuss WG adoption
 - Will need to arrange help from Security Directorate.

Motivations

Related Documents (Slide Two of Two)

- Rfc5662bis-00
 - Need to schedule WG adoption
- RFC8178
 - Need to reference as definitive in rfc5661bis
 - Already in -01.
- RFC8434
 - Need to incorporate in rfc5661bis
 - Probably post-02

Motivations

Quality issues found as we worked with docs

- Use of RFC2119 terms not in accord with the definition
- A lot of confused terminology to clean up:
 - Different meanings for “client owner”.
 - Handling of the word “verifier”
- A number of troubling areas
 - Handling of directory delegation needs clarification.
 - Requirements (e.g. atomicity) for persistent reply cache makes it unimplementable.
 - Discussion of memory-mapped IO does not make sense

What's in it

Finding out yourself

- For changes in -01, can just diff it against -00.
 - Most of what shows up is really a change, rather than an artifact
 - Should be explained in Appendix B.5.1
- For the changes, in -00 a diff against rfc8881 shows a lot of changes, but many are not real
 - Thousands of lines added/deleted/changed, but a lot are ignorable.
 - Missing page numbers in rfc8881
 - Reference format changes
 - Section number changes, due to reorganization of Sections 1-2.
 - Explanation in Appendices B.1 through B.3 or ask on list.

What's in it

Shift from rfc8881 to -00

- Organizational changes
 - New Section 1 with old Section 2 being reorganized.
 - Changes to adapt to internationalization, extension, and security being dealt with in separate documents
- Addressing a lot of errata reports
- Shift from “RECOMMENDED” to “recommended” attributes
- Addition of Appendices
 - **A**: Nature of changes
 - **B**: Status of changes
 - **C**: Issues Requiring Discussion

What's in it

Shift from -00 to -01 (Slide one of three)

- Revised handling of Retry/EOS (described in Appendix C.2.1)
 - Get rid of “MUST” meaning “It *really would be nice* if things worked this way”
 - Get rid of “MUST NOT” meaning “This is a bad idea ,and we feel *really strongly* about that”.
 - Makes explicit that EOS provides “at-most-once” semantics, which it does.
- Rewrote section about memory mapping and locking (described in Appendix C.2.3), because the old one
 - Neglected the effects of OPEN
 - Misunderstood byte-range locking. Needed correction since,
 - Advisory locking has no effect on IO.
 - Mandatory locking does *not* cause locks to be taken, leading to deadlock.

What's in it

Shift from -00 to -01 (Slide two of three)

- Terminology Clarifications
 - Distinguished “client owner” and “client owner id”
 - Explain multiple uses of the word “verifier”
 - pNFS terminology
 - Created “file data provider” to cover both “data server” and “storage device”.
 - Moved from ‘storage protocol” to “data protocol”
 - Clarified the discussion of “control protocol”

What's in it

Shift from -00 to -01 (Slide three of three)

- Re-organization of discussion of pNFS security
 - Security requirements are common to all mapping types
 - How those requirements are met differs for three mapping type groups
 - Data protocols using RPC with separate control protocol
 - Data protocols using RPC w/o separate control protocol
 - Data protocols not using RPC.
 - Threat analysis for pNFS will be organized similarly, but,
 - For data protocols using RFC, issues between client and data provider will be dealt with in the security document.
 - For control protocols, issues between server and file data provider are addressed in this document.

What's will be in it

Already scheduled for draft -02

- Remaining errata reports
- Revisit “RECOMMENDED” Attributes
 - Since are not all equally “recommended”
 - Particular issue for authorization-related attributes.
- Clean up description of Named Attributes.

What's will be in it

For draft -02 and beyond

- Items from Appendix C.2
 - Needed revisions for directory delegation (as described in C.2.2)
 - Needed work regarding persistence (as described in C.2.4)
 - Description of persistent reply cache presents it as unimplementable.
 - Possibility of lock persistence not clearly discussed
 - Reorganization of attribute descriptions
- Threat Analysis for pNFS
- Address remaining errata
- Other things that need doing

Topics for Discussion

Things done in -01

- Any problems with new text introduced by changes described in C.2.{1,3}?
- Any comments/issues regarding other, more minor changes?
- Any comments regarding work already scheduled for -02 (see Slide #11)
- Suggestions as to priorities for additional work in -02?

Topics for Discussion

Things mentioned in -01, but not yet done.

- Directory delegation
 - Clarify situations with regard to cookie changes.
 - Exclude batching and time-based delay from directory notifications
- Persistence Issues
 - Persistent sessions
 - Clarify atomicity requirements so they don't make implementation impossible
 - Allows server reboot to terminate COMPODs
 - Clarify possibility of lock persistence via continuation of existing clientid

Topics for Discussion

Anything else?

- Have I missed anything?
 - Almost certainly I have.
 - Please bring up issues as soon as possible.
 - If you want to discuss at a forthcoming meeting, best to raise them on list first.

Process Going Forward

My expectations as author

- Have the following channels for discussion:
 - Meetings such as IETF117 and this meeting.
 - Will discuss this document and future drafts when needed
 - At least monthly but not weekly
 - Need consensus regarding scheduling of Interims.
 - The WG mailing list
 - Best to discuss issues there first
 - Also a place for follow-up discussions
- Future Drafts
 - At least two before requesting WGLC