

NFSv4.1 – State of the Union

A specification update

Mike Eisler

Dave Noveck

Spencer Shepler

2008-05-12

Contents

- Changes since WG Last Call
- Changes left
- Next steps in IETF process
- After NFSv4.1
- Questions

Changes since WG Last Call

- Draft 19
- Draft 20
- Draft 21
- Draft 22
- Draft 23

Draft 19

- Massive keyword fix
 - e.g. mandatory->REQUIRED
- ADD device ID notifications dropped
- GETDEVICELIST reduced in scope to accommodate blocks and objects
 - ◆ Now returns just device IDs
 - ◆ No notifications
 - ◆ Only a single file system (fsid)

Draft 20

- Updated minor versioning rules to allow adding cases to switched unions
- Corrected data type names for Basic Types
- Corrected normative references
- Documented a scenario where TCP retransmit delays could result in loss of lease
- Added sections to discuss:
 - ♦ Stateid Seqid Values and Byte-range Locks
 - ♦ Issues with Multiple Open-owners
- Added `time_access` to list of attributes client can change under delegation
- Fixes to multi-server-name-space data types

Draft 21

- Server MAY re-compute highest slot ID values in replies to retries.
- Made a “session inactivity timer” an explicit part of protocol.
- Clarified that re-opens of the same file by the same open-owner bump the seqid of the stateid.
- Clarification of retention attributes
- NFS4ERR_WRONG_TYPE error replaced with NFS4ERR_BADTYPE in some cases
- Clarify when a client can provide a new offset (EOF) in LAYOUTCOMMIT (client must have had a _RW layout that includes the new offset).
- NFSv4.1 servers that are just data servers MUST get the lease period from the MDS that refers the the client to the data server.
- Made NFL4_UFLG_COMMIT_THRU_MDS a MUST not a preferred.
- Clarified that LAYOUTCOMMIT requires the client to have a _RW layout on the affected region.

Draft 22

- Servers SHOULD use the client ID of the session ID to determine whether SETATTR with a special stateid produces a recall.
- Further expanded on the role of TCP retransmissions in the loss of lease
- Clarified the use the “seqid” in layout stateids
 - Defined limits on number of parallel layout operations per file
- Deleted extraneous NFS4ERR_CONN_BINDING_NOT_ENFORCED error
- NFS4ERR_BAD_SESSION_DIGEST is just for SET_SSV
- Clarified current stateid
- Specified rules for handling of execute rights in the ACCESS operation
- Clarified spo_must_allow/ spo_must_enforce capabilities of EXCHANGE_ID and client ID
- LINK & RENAME allowed to return NFS4ERR_WRONGSEC
- CB_NULL MUST be implemented

Draft 23

- Added more items in the "Differences from NFSv4.0" section
- Clarified the caching of CREATE_SESSION replies
- Synchronized block and object layout recall types with base specification.
- Defined LAYOUTGET's loga_length field and specified formal rules for valid arguments and results
 - Formalized opportunistic LAYOUTGET attempts.

Next Steps in IETF Process

- Need to hand document to Transport Area Advisor at IETF
- Draft-23 might be that document
- Do not expect an RFC before the end of 2008

After NFSv4.1

- Parallel Access to metadata
- Security labeling (e.g. NSA proposal)
- De-Dup aware clients
- End-to-end data integrity
- ...

Do we really need new operations?

Do we have enough infrastructure to do NFSv4.next?

Questions?

thanks