

# NFSv4.1

## Current Status

Dave Noveck

February 5, 2007

# Outline

- NFSv4 Minor Versioning
- What's in Minor Version One?
  - Sessions
  - pNFS
  - Delegation improvements
  - Other stuff
  - Spec corrections and clarifications
- Status of Minor Version One

# NFSv4 Minor Versioning

- Allows Incremental Extensions
- Same Basic Protocol Structure
  - COMPOUND, etc.
- Mostly, optional extensions
  - New attributes, new operations in Compound
- Also, some required infrastructural change
  - Should be kept to a minimum

# Sessions

- Required Infrastructural Change
  - Group attempted as optional
    - Significant benefits of inclusion
    - Optionality made protocol too complicated
- Exactly-Once semantics
  - Dependable reply cache
  - Required size is bounded
  - Uses requester-assigned slots with an incrementing sequence within each

# Sessions, continued

- Better Callback support
  - Callbacks on same connection as requests
- Trunking
  - Multiple IP's can be identified as same server
  - Multiple connections in a single server
  - Multiple sessions for a single clientid
- Improves locking model
  - Every request has clientid, based on session
  - Owner-based sequencing gone

# Parallel NFS

- High-performance data access
  - Allows multiple servers to co-operate
- Separate meta-data and data access
  - Meta-data access via NFSv4.1
  - Data access may use a number of protocols
    - May use NFSv4.1 (described in v4.1 spec)
    - May use block access (FCP/iSCSI)
    - May use object Scsi extensions
    - Block and object access in separate specs

# pNFS, Layouts

- Layouts specify location of data
- Typically striped among multiple data servers
- Layout types for file, object, block access
- Recallable like delegations
  - Restriping of files
  - Server resource management
  - Conflicting uses (typically blocks)

# Delegation Improvements

- Client can request delegations back
- Delegation management improvements
  - Server can request reduction in count
  - Client returns its least useful layouts
  - Server has no valid way to pick those
  - Similarly for layouts
- Directory delegations
  - Only read directory delegations
  - Also notification mechanism for changes



# Multi-Server Namespace

- Replaces replication/migration chapter
- More detail on filesystem transitions
  - More details means more possibilities/choices
- Explains referrals explicitly
- `fs_locations_info` (`fs_locations++`)
  - Replica priorities
  - Information on fs transition implementation choices
  - Integrated with trunking
    - Some replicas may allow simultaneous use.

# ACL improvements

- Support for Windows dynamic inheritance
- Separate setting of security ACL's and auditing ACL's
- Much more explanation of
  - Interaction of mode and ACL's
  - Special ACE who-values like @OWNER
  - Interpretations of ACE mode bits

# Implementation ID's

- Identification of client and server
- Aid to support
  - Statistics on what you are interacting with
  - Bug tracking
- Not allowed to do anything different
  - Implementation may not provide ID
  - No protocol specials

# Retention Support

- Attributes to lock down files
  - No changes, removal, renaming
  - For compliance
- Can set attributes to:
  - Lock down file until a given time in the future
  - Lock down forever
  - Retain file based on specific events
  - Retain based on administrative holds (e.g. a court order)

# Additional Cleanup

- Uses xml2rfc
  - html-based spec with links for cross-refs
- Dot-x derivable from op descriptions
  - Avoids possibility of conflicts
- Corrects mistakes and contradictions
  - That stuff will also go into rfc3530bis
- Automated error tables
  - From ops to errors and vice versa

# Status

- Draft-08 is out (444 pages)
- Draft-09 out very soon
  - Complete protocol
  - No more .x changes
  - Unless testing shows up issues
- Future drafts
  - Still need more/better explanation
  - Clarifications/corrections
  - Results of interoperability testing

# Going Forward

- More implementation needed
  - Feature integration under way
    - e.g. pNFS plus Sessions
  - Want to have impl. experience before last call
- Need more detailed spec review
  - Lots of errors seen in RFC3530

# More Info

- Working group list: [nfsv4@ietf.org](mailto:nfsv4@ietf.org)
- V4 editor site: [www.nfsv4-editor.org](http://www.nfsv4-editor.org)
  - Issue tracker
  - Recent drafts including html versions
- [www.ietf.org](http://www.ietf.org)
  - Announcement of new drafts, including pNFS blocks and objects drafts