




Go further, faster®

IETF 84: NFSv4.2

Tom Haynes

thomas@netapp.com





Summary

- Weekly review of all content
 - About 12 meetings in all
- Wide representation from across the WG
- Editorial issues
 - Get rid of requirements and the sell job
- Protocol issues
 - Most features proposed by server developers
 - Need to get client developers engaged earlier



Server-side Copy

- Talk about authentication based rather than pull based
 - Editorial, no content changes
- Server reboot recovery and also lock recovery
 - Protocol changes
 - Add a new stateid to COPY and COPY_NOTIFY



IO_ADVISE

- Made it closer to `posix_fadvise`
 - Got rid of Direct I/O discussion
 - Do not over-constrain the server implementation
 - Do not talk about life time of hint
 - Do not talk about how many active hints
 - Keep them byte-range based, like `posix_fadvise`
 - Editorial change, kept the intent true



Space reservations

- `fallocate()` issue with Linux
 - set `space_reserved` attribute
 - set size of file
 - `fallocate()` needs offset as well



Holes

- INITIALIZE or WRITE_PLUS
 - COMMIT?
 - Asynchronous nature of operation
 - Use Server-side Copy tricks to call back to client
 - How to halt?
- ADBs rewritten as adb_holes
 - More of a hole than a block
- pNFS
 - Dumb DSes - respond with all local info in range
 - Smart clients - use layout info to weave it together
 - Example, 2 DSes, 1 Meg stripes, every other Meg is a hole



What next?

- Close on all AIs from review
 - Probably 1/2 done
 - Remaining has controversy
- Finish review
 - Server-side Copy network partition
 - Space reservations and fallocate()
 - Hole punching and asynchronous servers