

**Delegation of Stateids
IETF107 —
err, Interim Virtual Meeting
Apr 22, 2020**

Tom Haynes
loghyr@hammerspace.com

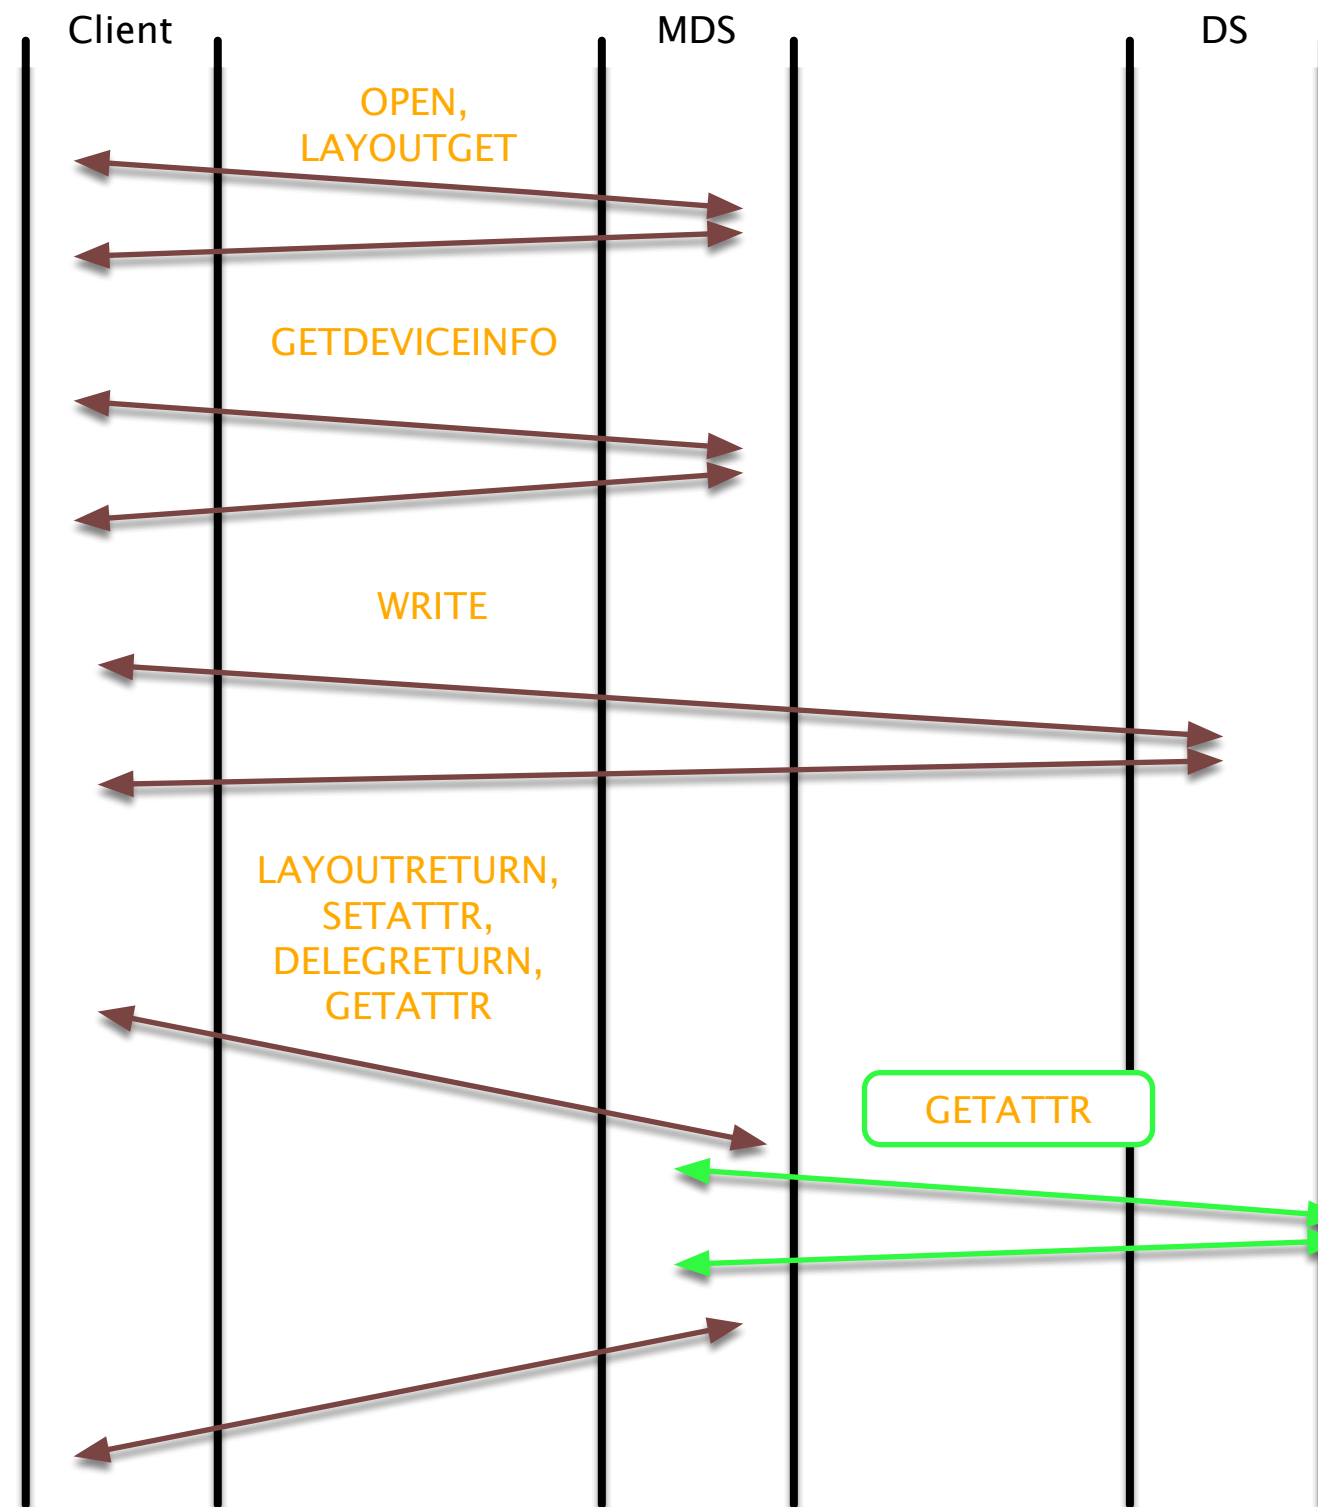
Agenda

- What we missed in IETF105?
- What is the current status?
- What will it take to push this work forward?

LAYOUT_WCC

- For a GETATTR(size/ctime/mtime), the MDS has to go to the DS
 - Client READs and WRITEs to the DS are invisible to the MDS
- Those READs and WRITEs return WCC to the client
- Client can feed that WCC to the MDS
- MDS can determine whether to use that WCC or not


echo “foo” > bar



Doesn't Attr Caching help?

- Client is just providing mtime and atime of metadata file
- MDS has to compare attrs of data file from last modification to now
 - If delta, then data file modified
- Even knowing metadata attrs changed, need to re-arm data file
- LAYOUT_WCC does this, eliminates GETATTR to DS

LAYOUT_WCC: echo “foo” > bar



LAYOUT_WCC vs LAYOUTRETURN vs LAYOUTCOMMIT

- LAYOUTRETURN mandates new version of RFC8435 - flexfiles layout type
- LAYOUTCOMMIT
 - Allows for layout type specific data
 - Strong cache consistency, not weak
 - Has additional semantics
- LAYOUT_WCC builds upon WCC in NFSv3 WRITE

The draft status

- Version 0 was everything but LAYOUT_WCC
- Agreed to be official WG item
 - Never finished off review items
- Still want next draft version to be an official WG item

Where we are?

- All but the WCC is implemented in a prototype form.
- Time skew is a killer
 - There is almost always a time adjustment when the delegation is returned
- Delegations hide sins from POSIX compliance checkers
 - E.g., create file, stat, add 1 byte, sleep 1, stat
 - All of that is done on the client and not the server

What is holding us up?

- Typical WG cycle is:
 - Get spec done
 - Implement code
 - Find we have incompatible implementations
- We want:
 - Get spec mostly done
 - Implement code
 - Get spec done
- Hung up on implementing code for WCC...