



**pNFS issues in NFSv4.1
Results from 2007-12-04
review**

**Mike Eisler
email2mre-ietf@yahoo.com**

- ▶ **Make deviceIDs 128 bits, opaque (no major/minor)**
- ▶ **When a deviceID to device address map changed, recall of layout NOT required**
- ▶ **Server to notify (via a new op) client when a deviceID has been changed**
 - Can give client option of gracefully draining in progress I/O
- ▶ **Recall layout by FSID/clientID and clientID will be supported**
- ▶ **Deleting a deviceID requires recalling all layouts that refer to that deviceID**
- ▶ **DeviceIDs now pertain to clientID and not FSID**
- ▶ **No stateIDs or delegations needed for deviceIDs or the entire deviceID map**

- ▶ **Client will be allowed to forget it has a particular range of a layout**
- ▶ **It must keep track of the layout stateid and the seqid field**
- ▶ **When server recalls a range the client has forgotten, it must return NFS4_OK**
 - **must follow up with a LAYOUTRETURN**
 - **This allows server and client to stay in sync on seqid in layout stateid**

- ▶ **The offset calculation in the spec does not work unless there is a single stripe pattern and unit size on the file**
- ▶ **DENSE packing seen as useful for less intelligent devices**
- ▶ **Anatoly and Suchit will provide simpler proposal than what was suggested in November**

- ▶ **Problem is that layout hint attribute (or any attribute) cannot be specified when file is created exclusively**
 - This motivates the recommendation to use persistent sessions
- ▶ **Resolution:**
 - Add EXCLUSIVE4_1 to createmode4 and add EXCLUSIVE4_1 case of createhow4 switched union
 - In addition to create verifier the new case has an fattr4
 - Add a new attribute that indicates what attributes are allowed in EXCLUSIVE4_1

- ▶ **We need more text to explain how the layout stateIDs seqid solves the race**

- ▶ **Agreed that access time not needed**
- ▶ **Still some remnants in spec that refer to `loca_time_access`**

- ▶ **LAYOUTCOMMIT vs. COMMIT**
- ▶ **Relationship of open mode and mode of layout needs description in LAYOUTGET**
- ▶ **Major/minor server owner and impact on pNFS**