

pNFS Block Layout: Layout Hint Issue

David L. Black, EMC
IETF nfsv4 WG

(pNFS) layout_hint attribute

- layout_hint is pNFS-layout-specific
- File and Object layouts
 - Tell server how to create file (e.g., striping)
- Block layout
 - Tell server what client's max I/O time is
- Problem discovered by implementers
 - Client max I/O time is per-client
 - layout_hint is specified as per-file
 - Setting this on every file is overkill for block layout

Fixing this problem

1. Decide whether it needs to be fixed (here).
 - Setting I/O time on each file is silly, at a minimum.
2. If yes, decide how to fix it (here).
3. Write new text to fix it:
 - Review on list
 - Send to RFC Editor via Area Director

Q1: Does this problem need to be fixed?

Possible Approaches

1. Tweak default behavior if attribute not set
 - Use most recent value provided by same client
2. Make attribute scope layout-dependent
 - File and object layouts: per-file
 - Block layout: per-filesystem (track per client)
 - Alternative: per-server
 - What if another client sets a high value?
3. Define a new attribute

Q2: Which approach?