

Parallel NFS (pNFS)
Flexible File Layout
NFSv4 WG
IETF99

Tom Haynes
loghyr@primarydata.com

ffds_stateid Conundrum

- Loosely coupled
 - NFSv3 no need
 - NFSv4.x has to be anonymous stateid
- Tightly coupled
 - Has to be a global stateid
- Feature or Bug?

Loosely Coupled Security

- Client and MDS connect to main KDC for metadata
- MDS acts as KDC for data
 - Gives client a ticket granting ticket to access storage device
 - Might want a smaller pool of synthetic ids
- Still fences by modifying the synthetic uid/gid

IANA Considerations

- pNFS Layout Types Registry
 - LAYOUT4_FLEX_FILES
- NFSv4 Recallable Object Types Registry
 - New
 - RCA4_TYPE_MASK_FF_LAYOUT_MIN
 - RCA4_TYPE_MASK_FF_LAYOUT_MAX
 - Old Missing
 - RCA4_TYPE_MASK_OTHER_LAYOUT_MIN
 - RCA4_TYPE_MASK_OTHER_LAYOUT_MAX

Write One Mirror

- New ff_flags4
 - FF_FLAGS_WRITE_ONE_MIRROR
- Client writes to one mirror
 - It updates all of the remaining
 - *Needs to report any errors to client*
 - COMMIT semantics defined

WGLC

- Version 6 has been in WGLC for a long time
- Currently at version 9 (really 11)
- Needs another review cycle
- Asked WG Chairs to reset the clock