

IETF90
NFSv4 WG
Flexible File Layout

Tom Haynes
Primary Data

Overview

- Started out as Object Layout v2
- Quickly morphed into Flexible File Layout
- Fit into the Requirements for Layout Types lingo
- Broke out what we want to implement now and what we might want later

Now versus Later: Pressure on new layout types

- Correct
 - Only one shot
- All encompassing
 - Get every possible new feature
- Future proof

Hasn't worked well for major revisions!

- 95% baked features get in
- Last 5% is killer
- We can't implement as is
- We can't get them out

Focus on what we can implement well

- Flexible file layout is about new coupling between non-tightly bound DS and MDS
- Striping and RAID can come later
- Build on top of an easy to implement Layout Type

Implementation experience

- Loosely coupled MDS and DS
- How does the DS communicate rsize and wrsize to the client?
 - Make MDS only use smallest of the possible DSes
 - Could change the GETDEVICEINFO response

GETDEVICEINFO response

```
55 struct ff_device_addr {  
56     multipath_list4 ffda_netaddrs;  
57     uint32_t        ffda_version;  
58     uint32_t        ffda_minorversion;  
59     uint32_t        ffda_rsize;  
60     uint32_t        ffda_wsize;  
61     bool            ffda_tightly_coupled;  
62 };
```

Status

- Tightly focused Layout Type
- Implementation experience
- Want to get this document as a WG document