



Go further, faster™

Metadata Striping in pNFS

draft-eisler-nfsv4-pnfs-metastripe-02.txt
IETF-79
2010-11-09

Mike Eisler





draft-eisler-nfsv4-pnfs-metastripe-02.txt

- Differences from last update
- NFSv4.x or standalone?
- Key points solved
- Issues of contention
- Data that supports proposal
- Overlaps
- Should proposal be combined/not combined?
- Who should review?
- Next steps



draft-eisler-nfsv4-pnfs-metastripe-02.txt

- Differences from last update
 - Added language for optimizing case of a file system that distributes entire directories as opposed to striping directories
 - Clarifications and bug fixes
 - No XDR changes
- NFSv4.x or standalone?
 - Standalone: 1 layout type
 - Requires NFSv4.1 or higher
- Key points solved
 - For NFS servers that distribute directory content and file metadata allows NFS clients to send request to optimal place in storage cluster



draft-eisler-nfsv4-pnfs-metastripe-02.txt

- Issues of contention
 - When directories are striped, optimizing LOOKUP, OPEN by name, etc. requires agreement on an algorithm for mapping {directory filehandle, “component name”} to a stripe index
 - First movers get the advantage but would need to license any IPR around their algorithms

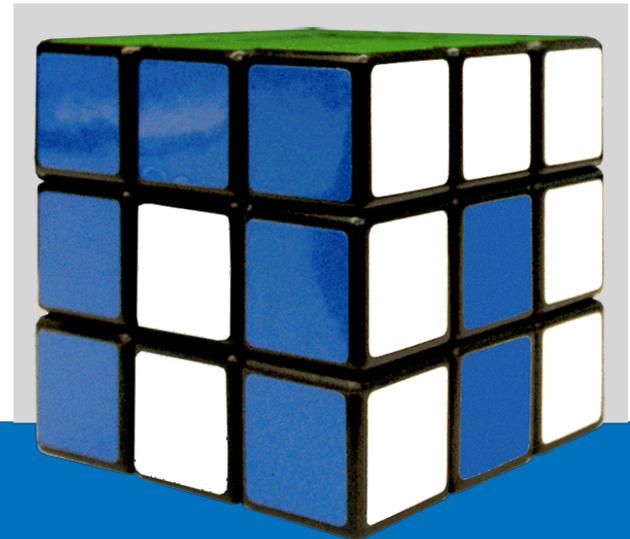


draft-eisler-nfsv4-pnfs-metastripe-02.txt

- Data that supports proposal
 - SPEC SFS 2008 benchmark is mostly operations other than READ or WRITE
 - Benchmark is derived from real customer data
 - Object storage and content repositories emerging
 - They often take the approach of eliminating the path name because of directory scaling concerns
- Overlaps
 - None
- Should proposal be combined/not combined?
 - N/A
- Who should review?
 - WG
- Next steps
 - Complete recovery pieces; doable by IETF 80
 - Prototyping from NetApp unlikely to occur there is client interest

Thanks

draft-eisler-nfsv4-pnfs-metastripe-02.txt



Questions and Answers