



Go further, faster™

Beyond NFSv4.1

Mike Eisler

mre DASH ietf AT eisler DOT com





Purpose of this Presentation

- Make the case for post NFSv4.1 standards work
- Start green lighting on work items



Case for post-NFSv4.1

- Meta-data intensive workloads
- Federations
- Virtualization



Is this all NFSv4.2?

- Ideally we should look to avoid adding more operations
- NFSv4 and NFSv4.1 have a lot of infrastructure to add new features:
 - attributes
 - layout types
 - unused flag fields



Meta-data intensive workloads

- E.g. workloads with intense LOOKUPS, READDIRs
- Delegations alone may not suffice
 - E.g. NFSv3 GETATTRs can overwhelm a single server
 - NFSv4.1 Delegation could also overwhelm a single device

Proposal: extend pNFS to support parallel MDS



Federations

- namespace
- Best practices and standards needed for
 - ID mapping
 - location of root of namespace
- Running code needed from clients to support migration
 - not just referrals



Virtualization

- NFS and Hypervisors make a great team today:
 - NFS is (practically) the most secure storage access protocol
 - files have more flexibility for provisioning than LUNs
- Hypervisors could use
 - hole punching
 - When a virtual machine (VM) de-allocates data indicate this to the NFS server
 - de-dupe awareness in the clients
 - Many VM images cloned from a master
 - Clones diverge but share a high %-age of content
 - Indicating which blocks are common will reduce overhead on hypervisors



Other Items that have been mentioned in the past ...

- End to End Data Integrity
- MAC Label and Enforcement
- Request Priority
- ...



Greenlighting on more things ...