

Controlling Storage Properties for NFSv4.2

Lakshmi Bairavasundaram -- NetApp

Peng Dai – Vmware

Pranoop Erasani -- NetApp

Christos Karamanolis – Vmware

David Noveck – EMC

Continuation of Work Presented at IETF78 in Maastricht

- *pNFS Storage Preference*
 - Presented by David Black of EMC
 - Co-authored by:
 - Mike Eisler of NetApp
 - Sorin Faibish of EMC
 - And others that were on both
- Changes here reflect WG comments then
 - Some, but not all
- In ***draft-dnoveck-storage-control-00***
 - *Forgot “-nfsv4-”. Sorry.*

People involved so far

Who	Comp.	Orig.	Draft	@IETF79
L. Bairavasundaram	NetApp	-	Y	-
D. Black	EMC	P	-	-
P. Dai	Vmware	Y	Y	Y
M. Eisler	NetApp	Y	-	Y
P. Erasani	NetApp	-	Y	Y
S. Faibish	EMC	Y	-	Y
C. Karamanolis	Vmware	Y	Y	-
D. Noveck	EMC	-	Y	P

Differences from IETF78 Talk

- XDR-based encoding
 - XML had been looked at (extensibility)
 - XDR worked fine although details of extensibility are going to require more care (like attribute structure).
- pNFS is now less central
 - Will talk about the reasons for this
 - Partly was the working group comments
- Specific set of properties
 - Need comments on completing that
- Future work to talk about
 - Need to decide on the minimum to be useful for v4.2.
 - Look, particularly, for things hard to do later
 - We know that some things will be deferrable
 - Complete set of properties would wait until v4.99 or beyond

Why are storage characteristics getting more important

- In other words, why do this?
 - Multiple reasons
- One issue is pNFS (data-metadata split)
 - Before that, storage could be controlled by putting the file in a given FS.
 - Now you have another variable (where the data is to be) with no easy means of control
- Also, the use of flash memory
 - Need to control what files get how much
- More different classes of things to store
 - Virtual devices for hypervisors

One-time situation or something ongoing?

- In other words,
 - Can we figure all the needed storage properties, write the draft, get it into v4.2, and then the issue is dealt with?
 - It doesn't look like it
- We are assuming storage properties will need their own extension mechanism
 - Also important to be able to prototype things in situations more like real use
 - “Working code” vs. “User-understandable property description”

Informative and Enforceable Properties

- Informative properties
 - Tells info useful in storage assignment
 - e.g. size, lifetime
- Enforceable properties
 - Specifies desired storage properties
 - Support for properties that are in the process of being complied with
 - Move one thing to make room for another
 - Two levels:
 - Advisory → Client wants those properties but can do w/o
 - Mandatory → Request failed in can't be secured

Getting to Ready-for-V4.2 State

- Cleanup and correction of what's there
 - Appreciate any comments you have
 - What is wrong. What is missing.
- What is an adequate set of properties?
 - Need to add some things to what is there
 - Appreciate property suggestions
 - With significant detail
 - Will have to accept imperfections, or debate forever
 - That's one of the reasons we need a property extension mechanism
- Two major additions (see next few slides)

Extension (1): Directory Handling

- Storage property inheritance suggested
 - In many cases, this will be a useful default
 - No real disagreement about it the fact that you need to do it.
 - But it does raise lots of issues for spec
 - CREATE-time or does RENAME have an effect?
 - What about LINK?
- Needs to be part of feature somehow
 - Details need to be worked out

Extension (2): More Compact Property Sets

- **Primarily for smaller files**
 - Naïve implementation of property sets could be big.
 - Could ask FS to share common property sets
 - Or could have the FS expose choices of property sets
- **Some proposals have been made in this area**
 - No time to decide issues and get them into a latest draft

Questions and Comments

- What is missing?
- What is not clear?
- What is there and shouldn't be?
- Are there additional people interested in participating in this work?
- Other questions?