

# NFS Version 4 Best Practices

Brian Pawlowski

# Practical status of NFS V4

- Several implementations emerging in production kernels
  - Linux, Solaris, NetApp, Humingbird, BSD, IBM, plus others (you know who you are)
- Q4'04 and Q1'05 will be interesting
- Reality is setting in for people - starting to sweat on testing and deployment strategies

# Particulars (not exhaustive)

- All clients can use flag at mount point (automount) to bind particular version - avoid V4
- NFS Version 4 in Linux is labelled EXPERIMENTAL
  - Not yet default mount option in 2.6.X
  - As an aside, TCP as NFS transport just came off EXPERIMENTAL
- On NetApp, server control via nfs.v4.enable
- Solaris has a min/max control that can be used to force allowed protocol to \*only\* V4 (or disable it)

# Implementation basics

- NFS version 4 is a client/server protocol
- NFS transparently negotiates highest version protocol - client binding driven
- Strongest security preferred
  - But sys admin can configure mount points for both “no security” and “Kerberos” - but to what end?
  - ACLs can be manipulated by V4 clients - are enforced for V2 and V3 clients (if the server supports it)
  - Fallout of V4 was Kerberos/GSS for V2 and V3 (layers)

# Interoperability issues

- V2 and V3 legacy interactions
- Worst case would be security reduction due to lack of strong GSS-based security
- Servers need to manage V2 and V3 clients regarding delegations and state disruptions
  - Need to document recommended behaviour (implementers)
  - Need to document resulting bad behaviour (degradation of V4 features by V3/V2 clients say) for deployment

# Resurrect the “implementation document”

- Focus on best practices
- Describe implementation recommendations on binding and recovery behaviour
- Recommended timer values
- Describe V4 behaviour in presence of V2/V3
- List implementations and their features (subject to change)

# Other use of implementations

- Document the protocol features implemented needed to advance standard

# Minor version items?

- Ability to renew a delegation from client
  - If reaped by server for operations like back-up (feedback from sys admins)
- NFS Version 3 “simulation” may best be done with separating lease timeouts from grace period for say lock recovery
- Need lots of input and feedback here!



# Proposal

- I have to write this document anyway - customers/vendors/your mom is getting worried
- Follow through with draft deliver by September 30
  - Intention to submit to WG for delivery as informational RFC
- ???

# Questions?