Time handling in NFS

Salit Gazit

helped by

Mike Eisler, David Robinson, Peter Staubach, Helen Chao and others...

Problem Definition

- NFS OTW time values are 32-bit unsigned, with an epoch of 1970.
- System time values are signed long.
- Issue: how do we handle/define time overflows:
  - 32-bit client or server.
  - 64-bit client or server.

When do we have a time overflow?

Whenever we are presented with a time value that cannot be represented as a positive 32-bit number either in the system time variable or in the OTW time variable.

Examples

- 0x80000000 in a 32-bit kernel - overflow.
- 0x80000000 in a 64-bit kernel - not overflow.
- -5 is always an overflow.
- 0x8000000000 in a 64-bit kernel - overflow.

So what?

- Negative time values are disallowed by the NFS protocol. We should not accept them.
- 64-bit clients and server can safely use the full 32-bit time values, thus allowing times beyond 2038.
- Reject values larger than 32-bit.

Possible Solutions

- Set pre-epoch times sent otw to zero, and large time values to INT32_MAX.
  - Problem: This is cheating...
- Return an error on time overflow. (*)
  - Problem: File will not accessible via NFS. Different behaviour in 32/64.

Issues

- Different max values for 32-bit and 64-bit kernels.
- Operations that contain invalid values in their arguments or return parameters should fail. Exception:
  - V3 ops with optional pre/post-op args can set them to FALSE.

- Invalid arguments should not be stored in the attribute cache.
- A 64-bit client may send a time value that will be rejected by a 32-bit server and vice versa.
- V2 OTW time variable is defined as an int32_t, even though the protocol requires uint32_t.

- UFS allows negative time values, but not times larger than INT32_MAX-1.
- A 32-bit program may behave differently when running on top of a 64-bit kernel or a 32-bit kernel.
  - Example: the stat system call, create system call.

- Different behaviour if same program runs on top of version 2 and version 3.
- Exposed bugs in overflow handling elsewhere in the system
  - Example: ufs, touch command.