
Subject: Re: [PATCH v2] SUNRPC: skip dead but not buried clients on PipeFS events

Posted by [Myklebust, Trond](#) on Wed, 25 Apr 2012 18:54:55 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2012-04-25 at 13:30 -0400, J. Bruce Fields wrote:

> On Fri, Apr 20, 2012 at 06:11:02PM +0400, Stanislav Kinsbursky wrote:

> > v2: atomic_inc_return() was replaced by atomic_inc_not_zero().

> >

> > These clients can't be safely dereferenced if their counter in 0.

>

> I'm pretty confused by how these notifiers work....

>

> rpc_release_client decrements cl_count to zero temporarily, to have it

> immediately re-incremented by rpc_free_auth.

>

> So if we're called concurrently with rpc_release_client then it's sort

> of random whether someone gets this callback.

>

> Is that a problem?

Not really. If we re-increment the client->cl_count in rpc_free_auth() then it would be so that we can send off a bunch of NULL rpc calls to destroy existing RPCSEC_GSS contexts. We shouldn't need to do any more upcalls in pipefs.

If we care, we could simply move the call to rpc_unregister_client() into rpc_free_auth() so that the pipefs notifier doesn't see us, or we could set a flag to have it ignore us.

> Also, is this an existing bug? (In which case Trond should take it
> now.)

--

Trond Myklebust

Linux NFS client maintainer

NetApp

Trond.Myklebust@netapp.com

www.netapp.com
