
Subject: Re: [PATCH v3] SUNRPC: set desired file system root before connecting local transports

Posted by [bfields](#) on Fri, 26 Oct 2012 17:52:54 GMT

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

On Wed, Oct 10, 2012 at 02:32:28PM +0400, Stanislav Kinsbursky wrote:

> >On Tue, Oct 09, 2012 at 03:47:42PM -0700, Eric W. Biederman wrote:

> >>"J. Bruce Fields" <bfields@fieldses.org> writes:

> >>

> >>>On Tue, Oct 09, 2012 at 01:20:48PM -0700, Eric W. Biederman wrote:

> >>>>"Myklebust, Trond" <Trond.Myklebust@netapp.com> writes:

> >>>>

> >>>>>On Tue, 2012-10-09 at 15:35 -0400, J. Bruce Fields wrote:

> >>>>>>Cc'ing Eric since I seem to recall he suggested doing it this way?

> >>>>>

> >>>>>Yes. On second look setting fs->root won't work. We need to change fs.

> >>>>>The problem is that by default all kernel threads share fs so changing

> >>>>>fs->root will have non-local consequences.

> >>>>

> >>>>Oh, huh. And we can't "unshare" it somehow?

> >>>

> >>>I don't fully understand how nfs uses kernel threads and work queues.

> >>>My general understanding is work queues reuse their kernel threads

> >>>between different users. So it is mostly a don't pollute your

> >>>environment thing. If there was a dedicated kernel thread for each

> >>>environment this would be trivial.

> >>>

> >>>What I was suggesting here is changing task->fs instead of

> >>>task->fs.root. That should just require task_lock().

> >>>

> >>>Oh, OK, got it--if that works, great.

> >>>

> >>>

> >>>The main problem with swapping fs struct is actually the same as in

> >>>root swapping. I.e. routines for copy fs_struct are not exported.

> >>>It could be done on place, but I don't think, that Al Viro would

> >>>support such implementation.

> >>>Trond?

It seems like we got stalled here.... Could you go ahead and try a patch, and see what people think?

--b.
