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

Posted by [bfields](#) on Wed, 10 Oct 2012 01:23:36 GMT

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

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.

> > Sorry, I don't know much about devtmpfs, are you suggesting it as a

> > model? What exactly should we look at?

>

> Roughly all I meant was that devtmpfsd is a kernel thread that runs

> with an unshared fs struct. Although I admit devtmpfsd is for all

> practical purposes a userspace daemon that just happens to run in kernel

> space.

Thanks for the explanation.

--b.
