
Subject: Re: [PATCH] nfs lockd reclaimer: Convert to kthread API

Posted by [akpm](#) on Thu, 19 Apr 2007 21:40:59 GMT

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

On Thu, 19 Apr 2007 17:19:24 -0400

Trond Myklebust <trond.myklebust@fys.uio.no> wrote:

> > Regardless kernel threads should be an implementation detail
> > not a part of the user interface. If kernel threads are part
> > of the user interface it makes them very hard to change.
> >
> > So it isn't that it doesn't make sense to me it is that it looks
> > fundamentally broken and like a maintenance nightmare.
> >
> > I would rather kill kernel threads then try and simulate them
> > when the kernel implementation has changed and kernel threads
> > are not visible.
> >
> > If I could be convinced that signal handling in kernel threads
> > is not something that will impede code modifications and refactoring
> > I would have less of a problem, and might not care.
>
> Tough. You're the one proposing to change existing code.

Using signals to communicate with kernel threads is fairly unpleasant, IMO.
We have much simpler, faster and more idiomatic ways of communicating
between threads in-kernel and there are better ways in which userspace can
communicate with the kernel - system calls, for example...

So I think generally any move which gets us away from using signals in
kernel threads is moving in a good direction.

> > With pid namespaces all kernel threads will disappear so how do
> > we cope with the problem when the sysadmin can not see the kernel
> > threads?
>
> Then you have a usability problem. How does the sysadmin reboot the
> system if there is no way to shut down the processes that are hanging on
> an unresponsive filesystem?

Where's the hang? A user process is stuck on h_rwsem?

If so, would it be appropriate to convert the user process to use
down_foo_interruptible(), so that the operator can just kill the user
process as expected, rather than having to futz around killing kernel
threads?

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
