Subject: Re: [PATCH] nfs lockd reclaimer: Convert to kthread API
Posted by ebiederm on Thu, 19 Apr 2007 19:20:29 GMT
View Forum Message <> Reply to Message

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

> On Thu, 2007-04-19 at 01:58 -0600, Eric W. Biederman wrote:
>> From: Eric W. Biederman <ebiederm@xmission.com>
>>
>> Start the reclaimer thread using kthread_run instead
>> of a combination of kernel_thread and daemonize.
>> The small amount of signal handling code is also removed
>> as it makes no sense and is a maintenance problem to handle
>> signals in kernel threads.
>
> Vetoed. Removing stuff just because it doesn't make sense to you is not
> acceptable.
>
> Signal handling in reclaimer threads is there in order to allow
> administrators to deal with the case where the server never comes up
> again.

Doesn't unmount handle that?

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.

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?

Eric

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