
Subject: Re: [PATCH] kthread: saa7134-tvaudio.c
Posted by [ebiederm](#) on Wed, 30 Aug 2006 12:39:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

ebiederm@xmission.com (Eric W. Biederman) writes:

> Andrew Morton <akpm@osdl.org> writes:
>
>> So in general, yes, the driver should be converted to the kthread API -
>> this is a requirement for virtualisation, but I forget why, and that's the
>> "standard" way of doing it.
>
> With the kthread api new kernel threads are started as children of keventd
> in well defined circumstances. If you don't do this kernel threads
> can wind up sharing weird parts of a parent process's resources and
> locking resources in the kernel long past the time when they are
> actually used by anything a user space process can kill.
>
> We have actually witnessed this problem with the kernels filesystem mount
> namespace. Mostly daemonize in the kernel unshares everything that
> could be a problem but the problem is sufficiently subtle it makes
> more sense to the change kernel threads. So these weird and subtle
> dependencies go away.
>
> So in essence the container work needs the new kthread api for the
> same reasons everyone else does it is just more pronounced in that
> case.

That plus the obvious bit. For the pid namespace we have to declare war on people storing a pid_t values. Either converting them to struct pid * or removing them entirely. Doing the kernel_thread to kthread conversion removes them entirely.

Eric

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
