
Subject: Re: [PATCH] saa7134-tvaudio: Convert to kthread API.
Posted by [Christoph Hellwig](#) on Fri, 20 Apr 2007 13:05:01 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, Apr 20, 2007 at 02:48:35PM +0200, Cedric Le Goater wrote:

> Eric W. Biederman wrote:

> > From: Eric W. Biederman <ebiederm@xmission.com> - unquoted

> >

> > It is my goal to replace all kernel code that handles signals
> > from user space, calls kernel_thread or calls daemonize. All
> > of which the kthread_api makes unnecessary. Handling signals
> > from user space is a maintenance problem because using a
> > kernel thread is an implementation detail and if user space
> > cares it does not allow us to change the implementation. Calling
> > daemonize is a problem because it has to undo a continually changing
> > set of state generated by user space, requiring the implementation
> > to change continually. kernel_thread is a problem because it
> > returns a pid_t value. Numeric pids are inherently racy and
> > in the presence of a pid namespace they are no longer global
> > making them useless for general use in the kernel.

> >

> > So this patch renames the pid member of struct saa7134_thread
> > started and changes its type from pid_t to int. All it
> > has ever been used for is to detect if the kernel thread
> > is has been started so this works.

> >

> > allow_signal(SIGTERM) and the calls to signal_pending have
> > been removed they are needed for the driver to operation.

> >

> > The startup of tvaudio_thread and tvaudio_thread_dep have
> > been modified to use kthread_run instead of a combination
> > of kernel_thread and daemonize.

> >

> > The result is code that is slightly simpler and more
> > maintainable.

>

> Here's a refreshed attempt using kthread_should_stop().

> Unfortunately, not tested bc we don't have the hardware.

I have a patch for this one flying around somewhere aswell.
It's trivial to kill the waitqueue and just use wake_up_process,
otherwise it looks pretty similar.

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