
Subject: Re: [PATCH] powerpc pseries eeh: Convert to kthread API

Posted by [linas](#) on Tue, 24 Apr 2007 17:24:48 GMT

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

On Tue, Apr 24, 2007 at 11:38:53AM +1000, Benjamin Herrenschmidt wrote:

> > The only reason for using threads here is to get the error recovery
> > out of an interrupt context (where errors may be detected), and then,
> > an hour later, decrement a counter (which is how we limit these to
> > 6 per hour). Thread reaping is "trivial", the thread just exits
> > after an hour.

>

> In addition, it should be a thread and not done from within keventd
> because :

>

> - It can take a long time (well, relatively but still too long for a
> work queue)

Uhh, 15 or 20 seconds even. That's a long time by any kernel standard.

> - The driver callbacks might need to use keventd or do flush_workqueue
> to synchronize with their own workqueues when doing an internal
> recovery.

>

> > Since these are events rare, I've no particular concern about
> > performance or resource consumption. The current code seems
> > to work just fine. :-)

>

> I think moving to kthread's is cleaner (just a wrapper around kernel
> threads that simplify dealing with reaping them out mostly) and I agree
> with Christoph that it would be nice to be able to "fire off" kthreads
> from interrupt context.. in many cases, we abuse work queues for things
> that should really done from kthreads instead (basically anything that
> takes more than a couple hundred microsecs or so).

It would be nice to have threads that can be "fired off" from an
interrupt context. That would simplify the EEH code slightly
(removing a few dozen lines of code that do this bounce).

I presume that various device drivers might find this useful as well.

--linas

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
