
Subject: Re: [PATCH] powerpc pseries eeh: Convert to kthread API
Posted by [ebiederm](#) on Tue, 24 Apr 2007 05:43:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

Benjamin Herrenschmidt <benh@kernel.crashing.org> writes:

>> Further in general it doesn't make sense to grab a module reference
>> and call that sufficient because we would like to request that the
>> module exits.
>
> Which is, btw, I think a total misdesign of our module stuff, but heh, I
> remember that lead to some flamewars back then...
>
> Like anything else, modules should have separated the entrypoints for
>
> - Initiating a removal request
> - Releasing the module
>
> The former is use did "rmmod", can unregister things from subsystems,
> etc... (and can file if the driver decides to refuse removal requests
> when it's busy doing things or whatever policy that module wants to
> implement).
>
> The later is called when all references to the modules have been
> dropped, it's a bit like the kref "release" (and could be implemented as
> one).
>
> If we had done that (simple) thing back then, module refcounting would
> have been much less of a problem... I remember some reasons why that was
> veto'ed but I didn't and still don't agree.

The basic point is because a thread can terminate sooner if we have an explicit request to stop, we need that in the design.

Because we need to find the threads to request that they stop we need to have some way to track them.

Since we need to have some way to track them having an explicit data structure that the callers manage seems to make sense.

Eric

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