Subject: Re: [PATCH 11/11] protect architectures where THREAD_SIZE >= PAGE_SIZE against fork bombs
Posted by David Rientjes on Tue, 26 Jun 2012 09:05:47 GMT
View Forum Message <> Reply to Message

On Tue, 26 Jun 2012, Glauber Costa wrote:

> > Right, because I'm sure that __GFP_KMEMCG will be used in additional
> > places outside of this patchset and it will be a shame if we have to
> > always add #ifdef's. I see no reason why we would care if __GFP_KMEMCG
> > was used when CONFIG_CGROUP_MEM_RES_CTLR_KMEM=n with the semantics that it
> > as in this patchset. It's much cleaner by making it 0x0 when disabled.
> >
>
> What I can do, instead, is to WARN_ON conditionally to the config option in
> the page allocator, and make sure no one is actually passing the flag in that
> case.
>

I don't think adding a conditional to the page allocator's fastpath when CONFIG_CGROUP_MEM_RES_CTLR_KMEM=n is appropriate. I don't understand why this can't be 0x0 for such a configuration, __GFP_KMEM certainly means nothing when we don't have it enabled so how is this different at all from kmemcheck?