
Subject: Re: checkpointing and restoring processes
Posted by [Dave Hansen](#) on Wed, 06 Jun 2007 15:27:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2007-06-06 at 13:37 +0200, Mark Pflueger wrote:

> hi everyone!

>

> i'm not subscribed to the list, so if you care to flame because of my noob

> question, just do it to the list, otherwise please cc me.

>

> i'm trying to write a checkpoint/restore module for processes and so have

> a basic version going already - problem is, when i restore the process,

> one of three things happens at random. first is, the process restored

> segfaults. second is, i get a kernel null pointer dereference and third

> is, i get a virtual address lookup error and a kernel crash. the trace

> back and the address always change.

Your patch definitely takes a simple, straightforward approach, which is good. But, there are a couple of things that need to get added.

For instance, when you make a copy of `tsk->mm`, what happens if that original task exits? It will drop its reference count and free that task, along with the `mm`. The new task will fault on its access to `newtsk->mm` because the `mm` has gone away.

Also, just setting `tsk->pid` is not enough to get the pid to show up in the system. It needs to make sure no other task has that pid as well as making entries in data structures like the pid allocation map.

In any case, it's nice to have other people interested in the same things! As Cedric suggested, please pop over to containers@lists.linux-foundation.org. There are at least two other efforts, besides ours working toward the same goal, so you'll have lots of comrades there. :)

-- Dave

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