
Subject: Re: Checkpoint/restart (was Re: [PATCH 0/4] - v2 - Object creation with a specified id)

Posted by [serue](#) on Tue, 22 Apr 2008 19:51:42 GMT

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

Quoting Dave Hansen (dave@linux.vnet.ibm.com):

> On Tue, 2008-04-22 at 23:36 +0400, Alexey Dobriyan wrote:

> > On Fri, Apr 18, 2008 at 07:44:59AM +0200, Nadia.Derbey@bull.net wrote:

> > > . echo "LONG XX" > /proc/self/task/<my_tid>/next_id

> > > next object to be created will have an id set to XX

> > > . echo "LONG<n> X0 ... X<n-1>" > /proc/self/task/<my_tid>/next_id

> > > next object to be created will have its ids set to XX0, ... X<n-1>

> > > This is particularly useful for processes that may have several ids if

> > > they belong to nested namespaces.

> >

> > Can we answer the following questions before merging this patch:

> >

> > a) should mainline kernel have checkpoint/restart feature at all

> > b) if yes, should it be done from kernel- or userspace?

> >

> > Until agreement will be "yes/from userspace" such patches don't make

> > sense in mainline.

>

> What do you mean by "from kernel" or "from userspace"? Userspace has to

> be involved somewhere, right? It at least need to open() the checkpoint

> file and pass it into the kernel.

Well let's start by answering a) :

I vote yes :-)

As for b), good question. Except as Dave points out I think it's pretty unlikely that we'll end up with "sys_checkpoint(checkpoint_dirnam)" and "sys_restart(checkpointed_dirnam)". Anything more user-space driven will likely require userspace to set ids on resources, right?

But yes, we have to completely answer b) before this patch goes in, no doubt about that.

Alexey, are you advocating for a completely kernel-driven approach, or just making sure we come to a decision?

thanks,
-serge

Containers mailing list

Containers@lists.linux-foundation.org

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