Subject: Re: [PATCH 6/7] containers (V7): BeanCounters over generic process containers
Posted by Paul Menage on Tue, 13 Feb 2007 09:03:12 GMT
View Forum Message <> Reply to Message

On 2/13/07, Pavel Emelianov <xemul@sw.ru> wrote:
> menage@google.com wrote:
> > This patch implements the BeanCounter resource control abstraction
> > over generic process containers. It contains the beancounter core
> > code, plus the numfiles resource counter. It doesn't currently contain
> > any of the memory tracking code or the code for switching beancounter
> > context in interrupts.
>
> Numfiles is not the most interesting place in beancounters.
> Kmemsize accounting is much more important actually.

Right, but the memory accouting was a much bigger and more intrusive patch than I wanted to include as an example.

>
> I have already pointed out the fact that this place
> will hurt performance too much. If we have some context
> on task this context must
> 1. be get-ed without any locking

Would you also be happy with the restriction that a task couldn't be moved in/out of a beancounter container by any task other than itself? If so, the beancounter can_attach() method could simply return false if current != tsk, and then you'd not need to worry about locking in this situation.

> 2. be settable to some temporary one without
> locking as well

I thought that we solved that problem by having a tmp_bc field in the task_struct that would take precedence over the main bc if it was non-null?

Paul