Subject: Re: [RFC] Containers infrastructure problems
Posted by xemul on Wed, 07 Mar 2007 07:05:13 GMT
View Forum Message <> Reply to Message

Paul Menage wrote:
> On 3/5/07, Pavel Emelianov <xemul@sw.ru> wrote:
>> Hi.
>>
>> I'm trying to implement RSS accounting via containers and
>> I have some difficulties and proposals.
>>
>> 1. Fork
>>
>> container_fork() is placed before new task obtains its
>> new mm_struct, files_struct, signal_struct etc. Isn't it
>> better to move container fork at the place where newly
>> created task it fully initialized to give controller
>> possibility to work with new mm, signals etc?
>
> Yes, that seems reasonable.
>
>>
>> 2. Early container usage
>>
>> Consider the following code:
>>
>>   struct my_container *cnt;
>>
>>   cnt = my_cnt_from_cont(task_container(current, &my_subsys));
>>
>> the problem is that when it is used before I register my
>> rss subsystem in initcall task_container returns me
>> dummytop container which is not my_container actually :(
>
> By definition all tasks are in dummytop (the top container in the
> dummy hierarchy) since you can't create sub-containers in the dummy
> hierarchy.
>
> You're right that before you're registered, the current container
> won't have a pointer for your subsystem. But calling
> container_register_subsys() from container_rss_init_early(), and
> having that called early on from init/main.c should be OK.

It is OK, but ->create callback should be aware of
the fact it is called on system boot time and thus
it mustn't call kmalloc() :)

>

> Paul
>