

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 is 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.

Paul

---