
Subject: Re: [ckrm-tech] [PATCH 7/7] containers (V7): Container interface to nsproxy subsystem

Posted by [xemul](#) on Thu, 05 Apr 2007 07:18:59 GMT

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

Paul Menage wrote:

> On 4/4/07, Pavel Emelianov <xemul@sw.ru> wrote:

>> Paul Menage wrote:

>> > On 4/3/07, Kirill Korotaev <dev@sw.ru> wrote:

>> >>

>> >> Sounds reasonable.

>> >> Pavel is preparing new RSS patches on top of your patches

>> >> which take into account other comments and Andrew objections.

>> >> Can we cooperate to send it altogether then?

>> >

>> > Sure. Do you want to send me any core container changes that you have?

>>

>> Actually I don't have any right now. Everything is built at

>> the top of 2.6.20 patches you sent.

>>

>> Maybe you will send me (give an URL) the preliminary patches

>> for more fresh kernel and I'll check what is still needed

>> for RSS container?

>>

>> I remind that I had only two problems with containers:

>> 1. fork hook was called to early - before the new task

>> initializes completely

>

> OK, looking at this, it's not clear where the best place to call it

> is. "As late as possible" sounds like a reasonable idea, but that

> messes up the current support for the nsproxy container subsystem,

> which wants to be able to move the current task into a new container

> based on nsproxy unsharing, after we've added the new task to its

> parent. I can imagine other subsystems maybe wanting to clone the

> current container at fork time too.

>

> That sort of implies that we need to split the container fork

> mechanism up into two parts, one early to add the refcount to the

> parent's container_group, and one late to handle the callbacks if

> desired. But that should be pretty straightforward.

Splitting sounds good. I'll try to prepare an appropriate patch.

>> 2. early need for rss containers (earlier than initcalls

>> are called)

>

> Couldn't you copy the way cpuset_init_early() works?

I did this in my previous version of rss container,
but I think it can be generalized. This is initialization
code that must not be nice-looking, but if it can be it
should be.

> Paul
>

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