
Subject: Re: [RFC][PATCH 1/5] Virtualization/containers: startup
Posted by [Hubertus Franke](#) on Fri, 03 Feb 2006 20:34:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

Greg KH wrote:

> On Fri, Feb 03, 2006 at 10:34:01AM -0800, Dave Hansen wrote:
>
>> Lastly, is this a place for krefs? I don't see a real need for a
>> destructor yet, but the idea is fresh in my mind.
>
>
> Well, what happens when you drop the last reference to this container?
> Right now, your patch doesn't cause anything to happen, and if that's
> acceptable, then fine, you don't need to use a struct kref.
>
> But if not, then why have a reference count at all? :)
>
> thanks,
>
> greg k-h
>

Greg ...

In our pid virtualization patches we removed the object automatically.
I presume that Kirill does the same although the code for that is not released.

I am working on converting our pid stuff so to follow the recommendation of Linux/Alan
to just do <container,pid> isolation through pidspaces.

I could release an allocation/ before hand. Based on Kirill's 1/5 patch adoption
(which took Linus's & Dave's naming comments into account)

How do we want to create the container?
In our patch we did it through a /proc/container filesystem.
Which created the container object and then on fork/exec switched over.

How about an additional sys_exec_container(exec_args + "container_name").
This does all the work like exec, but creates new container
with name "..." and attaches task to new container.
If name exists, an error -EEXIST will be raised !

-- Hubertus
