
Subject: Re: [RFC][PATCH 1/5] Virtualization/containers: startup

Posted by [dev](#) on Sun, 05 Feb 2006 15:05:44 GMT

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

> I just did a global s/vps/container/ and it looks pretty reasonable, at
> least from my point of view.

>

> Couple of minor naming nitpick questions, though. Is vps/container_info
> really what we want to call it? It seems to me to be the basis for a
> real "container", without the _info part.

Can be omitted.

> "tsk->owner_container" That makes it sound like a pointer to the "task
> owner's container". How about "owning_container"? The "container
> owning this task". Or, maybe just "container"?

This is why I don't like "container" name.

Please, also note, in OpenVZ we have 2 pointers on task_struct:

One is owner of a task (owner_env), 2nd is a current context (exec_env).
exec_env pointer is used to avoid adding of additional argument to all
the functions where current context is required.

Linus, does such approach makes sense to you or you prefer us to add
additional args to functions where container pointer is needed? This
looks undersirable for embedded guys and increases stack usage/code size
when virtualization is off, which doesn't look good for me.

> Any particular reason for the "u32 id" in the vps_info struct as opposed
> to one of the more generic types? Do we want to abstract this one in
> the same way we do pid_t?

VPS ID is passed to/from user space APIs and when you have a cluster
with different archs and VPSs it is better to have something in common
for managing this.

> The "host" in "host_container_info" doesn't mean much to me. Though, I
> guess it has some context in the UML space. Would "init_container_info"
> or "root_container_info" be more descriptive?

init_container?

(Again, this is why container doesn't sound for me :))

> 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.

I don't see much need for krefs, do you?

In OpenVZ we have 2-level refcounting (mentioned recently by Linus as in
mm). Process counter is used to decide when container should
collapse/cleaned up and real refcounter is used to free the structures
which can be referenced from somewhere else.

Also there is some probability that use of krefs will result in sysfs :)))

> How does the attached patch look?

I will resend all 5 patches tomorrow. And maybe more.

Kirill
