
Subject: Re: [RFC][PATCH 1/5] Virtualization/containers: startup
Posted by [Linus Torvalds](#) on Mon, 06 Feb 2006 16:56:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Sun, 5 Feb 2006, Kirill Korotaev wrote:

>
> 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.

That naming _has_ to change.

"exec" has a very clear meaning in unix: it talks about the notion of switching to another process image, or perhaps the bit that says that a file contains an image that can be executed. It has nothing to do with "current".

What you seem to be talking about is the _effective_ environment. Ie the same way we have "uid" and "euid", you'd have a "container" and the "effective container".

The "owner" name also makes no sense. The security context doesn't "own" tasks. A task is _part_ of a context.

So if some people don't like "container", how about just calling it "context"? The downside of that name is that it's very commonly used in the kernel, because a lot of things have "contexts". That's why "container" would be a lot better.

I'd suggest

current->container - the current EFFECTIVE container
current->master_container - the "long term" container.

(replace "master" with some other non-S&M term if you want)

(It would make sense to just have the prepend-"e" semantics of uid/gid, but the fact is, "euid/egid" has a long unix history and is readable only for that reason. The same wouldn't be true of containers. And "effective_container" is probably too long to use for the field that is actually the _common_ case. Thus the above suggestion).

Linus
