
Subject: Re: [PATCH 1/4] Virtualization/containers: introduction

Posted by [ebiederm](#) on Tue, 07 Feb 2006 03:34:33 GMT

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

Kirill Korotaev <dev@openvz.org> writes:

> Hello,

>

> I tried to take into account all the comments from you guys (thanks a lot for
> them!) and prepared a new version of virtualization patches. I will send only 4
> patches today, just not to overflow everyone and keep it clear/tidy/possible to
> review.

>

> This patch introduces some abstract container kernel structure and a number of
> operations on it.

>

> The important properties of the proposed container implementation:

> - each container has unique ID in the system

What namespace does this ID live in?

> - each process in the kernel can belong to one container only
Reasonable.

> - effective container pointer (econtainer()) is used on the task to avoid
> insertion of additional argument "container" to all functions where it is
> required.

Why is that desirable?

> - kernel compilation with disabled virtualization should result in old good
> linux kernel
A reasonable goal.

Why do we need a container structure to hold pointers to other pointers?

> Patches following this one will be used for virtualization of the kernel
> resources based on this container infrastructure, including those VPID patches I
> sent before. Every virtualized resource can be given separate config option if
> needed (just give me to know if it is desired).

>

> Signed-Off-By: Kirill Korotaev <dev@openvz.org>

>

> Kirill

>

> P.S. I understand that this virtualization spam can be uninteresting for some of
> you, just give me to know if you want to be removed from CC.

May I please be added to the CC list.

We are never going to form a consensus if all of the people doing implementations don't talk.

Eric
