
Subject: Re: [RFC][PATCH 2/7] VPIDs: pid/vpid conversions
Posted by [Alexey Kuznetsov](#) on Thu, 09 Feb 2006 01:11:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello!

> In capability.c it does for_each_thread or something like that. It is
> very similar to cap_set_pg. But in a virtual context all != all :)

Do you mean that VPID patch does not include this? Absolutely.
VPIDs are not to limit access, the patch virtualizes pids, rather
than deals with access policy.

Take the whole openvz. Make patch -R < vpid_patch. The result is perfectly
working openvz. Only pids are not virtual, which does not matter. Capisco?

> I think for people doing migration a private pid space in some form is
> necessary,

Not "private", but "virtual". VPIDs are made only for migration, not for fun.

And word "private" is critical, f.e. for us preserving some form of pid
space is critical. It is very sad, but we cannot do anything with this,
customers will not allow to change status quo.

> My problem with the vpid case and it's translate at the kernel
> boundary is that boundary is huge

Believe me, it is surprizingly small.

Alexey
