
Subject: Re: [RFC][PATCH 5/7] VPIDs: vpid/pid conversion in VPID enabled case
Posted by [Alexey Kuznetsov](#) on Fri, 03 Feb 2006 10:52:02 GMT

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

Hello!

> This is an interesting approach. Could you elaborate a bit on why
> you need the two different approaches? What conditions cause the switch
> to the sparse approach?
>
> Also, if you could separate those two approaches out into two different
> patches, it would be much easier to get a grasp about what's going on.
> One of them is just an optimization, right?

Exactly. They are not two different "approaches", it is just a simple optimization.

In our approach each process has pair of pids: one is global unique identifier (read, it is real pid from viewpoint of kernel), another is virtual pid. We just do not want to lose cycles allocating both of them all the time, so we derive virtual pid from global one, it is automatically "virtually" unique.

Switch to "sparse" mapping happens only after migration, when a process must be recreated with the same virtual pid, but with new global pid.

> Did you happen to catch Linus's mail about his preferred approach?
>
> <http://marc.theaimsgroup.com/?l=linux-kernel&m=113874154731279&w=2>

Of course. Logically, it would be final solution.

VPID approach is pragmatic: it does not modify existing logic, rather it relies on it. So, it just allows to use virtual pids in a simple and efficient way, which is enough for all known tasks.

Alexey
