
Subject: Re: [PATCH 0/13] Pid namespaces (OpenVZ view)

Posted by [ebiederm](#) on Tue, 29 May 2007 04:30:45 GMT

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

Pavel Emelianov <xemul@openvz.org> writes:

> Hmm. I see. So you don't care that the pids in the namespace #2 are still
> the same. I can understand that politics for namespace #1, but for #2...

I'm confused, I think the statement above is wrong.

If we just checkpoint/restart a leaf pid namespace we don't care about
the other pids, in other namespace.

If we checkpoint/restart a pid namespace with another pid namespace
nested inside it we need to preserve the pids in the pid namespace we
are checkpointing and in a nested pid namespaces.

Pids in namespaces that none of the process we are migrating cannot
see we do not care about. (i.e. the init pid namespace, and possibly
some of it's children)

> OK, if you need this let us go on with such model, but I'd like to see
> the CONFIG_PID_NS_MULTILEVEL for this. Or at least CONFIG_PID_NS_FLAT for
> my model as we do not need to sacrifice the performance to such generic
> behavior.

Where is the world would a performance sacrifice come in? If you
happen to be using a deeply nested pid namespace I can see a small
performance hit, there is fundamentally more to do. However if you
don't use a nested pid namespace there should not be more work todo
and it should be impossible to measure the over head.

Further 3 levels should be as simple to implement and as cheap as two
levels. Because we can continue to use static allocation.

Eric
