
Subject: Re: [PATCH 0/28] Pid namespaces (two models)

Posted by [akpm](#) on Tue, 19 Jun 2007 23:00:09 GMT

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

On Fri, 15 Jun 2007 19:55:43 +0400

Pavel Emelianov <xemul@openvz.org> wrote:

> Long ago Sukadev and I sent two approaches for pid namespaces - the
> hierarchical model in which namespaces are nested into each other,
> and the flat model, where pids have only two values and creation of
> level 3 namespace is prohibited.

>

> After that I showed that multilevel model introduces a noticeable
> overhead of approximately 1-2% to kernel standard operations like
> fork() and getpid(). At the same time flat model showed no performance
> hit on these tests.

>

> Nevertheless multilevel model is worth living.

>

> This set introduces booth models each under its config option. The
> set is logically splitted into the following parts:

Making this configurable sounds like a very bad idea to me, from the
maintainablility/testability/understandability POV.

We should just make up our minds and do it one way, do it right?

I assume that means hierarchical.

> The following tests were run:

> [1] nptl perf test

> [2] getpid() speed

> [3] ltp (not for speed, but for kernel API checks)

>

> The testing results summary:

> * Flat model provides zero overhead in init namespace for all the
> tests and less than 7% in the namespace for nptl test only.

> * Multilevel model provides up to 2% overhead in init namespace and
> more than 10% for nptl test in the level 2 namespace.

>

So that means we take a 3% hit in these operations when PID_NS_MULTILEVEL
is enabled but the system isn't using containers at all?

If so, I'm surprised that the cost is this high. This should be the first
thing we should optimise and I bet there's some quicky way of doing it.

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
