
Subject: Re: Pid namespaces approaches testing results
Posted by [Pavel Emelianov](#) on Wed, 06 Jun 2007 07:09:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

Cedric Le Goater wrote:

>> The flat model has many optimization ways in comparison with the multilevel
>> one. Like we can cache the pid value on structs and some other.
>>
>> Moreover having generic level nesting sounds reasonable. Having single level
>> nesting - too as all the namespace we have are single nested. But having the
>> 4 level nesting sounds strange... Why 4? Why not 5? What if I don't know how
>> many I will need exactly, but do know that it will be definitely more than 1?
>>
>> Moreover - I have shown that we can have 1% or less performance on generic
>> nesting model, why not keep it?
>
> did you send that patchset ? is it included in the one you sent ?

The patchset I sent earlier changed slightly. The tests were performed on the version I sent. Right now I'm waiting for your results to make a final decision whether or not to develop the flat model together with the hierarchical one.

So what are we going to do? The ways we have:

1. Make two models - hierarchical and flat. Maybe we'll see how to merge them later;
2. Optimize the hierarchical model to produce no performance hit on the first 2 levels (init and VS). I don't see the way to make this gracefully, but I maybe this can be solved ... somehow. Anyway, if the latest patches from Suka do not produce any noticeable overhead, I am OK to go on with them;
3. Make the CONFIG_MAX_NS_DEPTH model. This is likely to be fast in the flat case, but I am in doubt whether Andrew will like it :)

> sorry if i missed something :(
>
> C.
>

Thanks,
Pavel

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