
Subject: Re: [PATCH 0/9] namespaces: Introduction
Posted by [ebiederm](#) on Fri, 19 May 2006 16:27:32 GMT
[View Forum Message](#) <> [Reply to Message](#)

Andrew Morton <akpm@osdl.org> writes:

> Herbert Poetzl <herbert@13thfloor.at> wrote:
>>
>> let me
>> give a simple example here:
>
> Examples are useful.
>
>> "pid virtualization"
>>
>> - Linux-VServer doesn't really need that right now.
>> we are perfectly fine with "pid isolation" here, we
>> only "virtualize" the init pid to make pstree happy
>>
>> - Snapshot/Restart and Migration will require "full"
>> pid virtualization (that's where Eric and OpenVZ
>> are heading towards)
>
> snapshot/restart/migration worry me. If they require complete
> serialisation of complex kernel data structures then we have a problem,
> because it means that any time anyone changes such a structure they need to
> update (and test) the serialisation.

There is a strict limit to what is user visible, and if it isn't user visible
we will never need it in a checkpoint. So internal implementation details
should not matter.

> This may be a show-stopper, in which case maybe we only need to virtualise
> pid #1.

Except we do need something for pid isolation, and a pid namespace is
quite possibly the light weight solution. If you can't see the pid it is
clearly isolated from you.

> Anyway. Thanks, guys. It sound like most of this work will be nicely
> separable so we can think about each bit as it comes along.

Yes, and there are enough issues it is significant.

Eric
