
Subject: Re: [PATCH 0/9] namespaces: Introduction
Posted by [Andrew Morton](#) on Fri, 19 May 2006 16:40:47 GMT
[View Forum Message](#) <> [Reply to Message](#)

ebiederm@xmission.com (Eric W. Biederman) wrote:

>
> > 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.

Migration of currently-open sockets (for example) would require storing of
a lot of state, wouldn't it?
