
Subject: Re: [PATCH 0/9] namespaces: Introduction
Posted by [Sam Vilain](#) on Sat, 20 May 2006 00:16:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Fri, 2006-05-19 at 05:41 -0600, Eric W. Biederman wrote:

> > It would help set minds at ease if someone could produce a
> > bullet-point list of what features the kernel will need to get it to the
> > stage where "most or all vserver and openvz functionality can be
> > implemented by controlling resource namespaces from userspace." Then we
> > can discuss that list, make sure that everyone's pretty much in
> > agreement.
> So this is slightly the wrong question. If you look at Sam's list you

Yes - the wrong question because it's too top down and encourages hacks :) It's wrong for the purposes of planning an implementation, but ok for easing minds about what will be covered, I think.

> will see that there are several independent dimensions to the complete
> solution. Most of them dealing with the increase in the number of users
> and the amount of work that is happening on a single kernel in this
> context.
>
> Basically we need to expect a lot of kernel tuning after we get the
> basics working.
>
> The proper question is: What needs to happen before we can run separate
> user space instances?

My guess would be most of the points under "isolation". The others are really just fine tuning / resource partitioning and fixing various things that break under virtualisation because of their design (eg, quota).

Sam.
