Subject: Re: Re: [RFC] network namespaces
Posted by kir on Wed, 06 Sep 2006 17:36:17 GMT
View Forum Message <> Reply to Message

Herbert Poetzl wrote:
> my point (until we have an implementation which clearly
> shows that performance is equal/better to isolation)
> is simply this:
>
> of course, you can 'simulate' or 'construct' all the
> isolation scenarios with kernel bridging and routing
> and tricky injection/marking of packets, but, this
> usually comes with an overhead ...
>
Well, TANSTAAFL*, and pretty much everything comes with an overhead.
Multitasking comes with the (scheduler, context switch, CPU cache, etc.)
overhead -- is that the reason to abandon it? OpenVZ and Linux-VServer
resource management also adds some overhead -- do we want to throw it away?

The question is not just "equal or better performance", the question is
"what do we get and how much we pay for it".

Finally, as I understand both network isolation and network
virtualization (both level2 and level3) can happily co-exist. We do have
several filesystems in kernel. Let's have several network virtualization
approaches, and let a user choose. Is that makes sense?

* -- http://en.wikipedia.org/wiki/TANSTAAFL