Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by Herbert Poetzl on Tue, 27 Jun 2006 16:09:08 GMT
View Forum Message <> Reply to Message

On Tue, Jun 27, 2006 at 01:54:51PM +0400, Kirill Korotaev wrote:
> >>My point is that if you make namespace tagging at routing time, and
> >>your packets are being routed only once, you lose the ability
> >>to have separate routing tables in each namespace.
> >
> >
> >Right. What is the advantage of having separate the routing tables ?

> it is impossible to have bridged networking, tun/tap and many other
> features without it. I even doubt that it is possible to introduce
> private netfilter rules w/o virtualization of routing.

why? iptables work quite fine on a typical linux
system when you 'delegate' certain functionality
to certain chains (i.e. doesn't require access to
all of them)

> The question is do we want to have fully featured namespaces which
> allow to create isolated virtual environments with semantics and
> behaviour of standalone linux box or do we want to introduce some
> hacks with new rules/restrictions to meet ones goals only?

well, soemtimes 'hacks' are not only simpler but also
a much better solution for a given problem than the
straight forward approach ...

for example, you won't have multiple routing tables
in a kernel where this feature is disabled, no?
so why should it affect a guest, or require modified
apps inside a guest when we would decide to provide
only a single routing table?

> From my POV, fully virtualized namespaces are the future.

the future is already there, it's called Xen or UML, or QEMU :)

> It is what makes virtualization solution usable (w/o apps
> modifications), provides all the features and doesn't require much
> efforts from people to be used.

and what if they want to use virtualization inside
their guests? where do you draw the line?

best,

Herbert

> Thanks,
> Kirill