Subject: Re: Re: [RFC] network namespaces
Posted by Herbert Poetzl on Fri, 08 Sep 2006 18:11:54 GMT
View Forum Message <> Reply to Message

On Fri, Sep 08, 2006 at 05:10:08PM +0400, Dmitry Mishin wrote:
> On Thursday 07 September 2006 21:27, Herbert Poetzl wrote:
> > well, who said that you need to have things like RAW sockets
> > or other protocols except IP, not to speak of iptable and
> > routing entries ...
> >
> > folks who _want_ full network virtualization can use the
> > more complete virtual setup and be happy ...

> Let's think about how to implement this.

> As I understood VServer's design, your proposal is to split
> CAP_NET_ADMIN to multiple capabilities and use them if required. So,
> for your light-weight container it is enough to implement context
> isolation for protected by CAP_NET_IP capability (for example) code
> and put 'if (!capable(CAP_NET_*))' checks to all other places.

actually the light-weight ip isolation runs perfectly
fine _without_ CAP_NET_ADMIN, as you do not want the
guest to be able to mess with the 'configured' ips at
all (not to speak of interfaces here)

best,
Herbert

> But this could be easily implemented over OpenVZ code by
> CAP_VE_NET_ADMIN split.
>
> So, the question is:
> Could you point out the places in Andrey's implementation of network
> namespaces, which prevents you to add CAP_NET_ADMIN separation later?
>
> --
> Thanks,
> Dmitry.