
Subject: Re: Re: [RFC] network namespaces

Posted by [Herbert Poetzl](#) on Mon, 11 Sep 2006 14:57:24 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Mon, Sep 11, 2006 at 04:40:59PM +0200, Daniel Lezcano wrote:

> Dmitry Mishin wrote:

> > On Friday 08 September 2006 22:11, Herbert Poetzl wrote:

> >

> >> 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)

> >

> > It was only an example. I'm thinking about how to implement flexible

> > solution, which permits light-weight ip isolation as well as full-fledged

> > network virtualization. Another solution is to split CONFIG_NET_NAMESPACE.

> > Is it good for you?

>

> Hi Dmitry,

>

> I am currently working on this and I am finishing a prototype bringing

> isolation at the ip layer. The prototype code is very closed to

> Andrey's patches at TCP/UDP level. So the next step is to merge the

> prototype code with the existing network namespace layer 2 isolation.

you might want to take a look at the current Linux-VServer
implementation for the network isolation too, should be
quite similar to Andrey's approach, but maybe you can
gather some additional information from there

> IHMO, the solution of splitting CONFIG_NET_NS into CONFIG_L2_NET_NS

> and CONFIG_L3_NET_NS is for me not acceptable because you will need

> to recompile the kernel. The proper way is certainly to have a

> specific flag for the unshare, something like CLONE_NEW_L2_NET and

> CLONE_NEW_L3_NET for example.

I completely agree here, we need a separate namespace
for that, so that we can combine isolation and virtualization
as needed, unless the bind restrictions can be completely
expressed with an additional mangle or filter table (as
was suggested)

best,
Herbert

> -- Daniel

Containers mailing list

