
Subject: Re: Re: [RFC] network namespaces
Posted by [ebiederm](#) on Sun, 10 Sep 2006 03:41:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

Herbert Poetzl <herbert@13thfloor.at> writes:

> On Sat, Sep 09, 2006 at 11:57:24AM +0400, 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?
>
> well, I think it would be best to have both, as
> they are complementary to some degree, and IMHO
> both, the full virtualization _and_ the isolation
> will require a separate namespace to work, I also
> think that limiting the isolation to something
> very simple (like one IP + network or so) would
> be acceptable for a start, because especially
> multi IP or network range checks require a little
> more effort to get them right ...
>
> I do not think that folks would want to recompile
> their kernel just to get a light-weight guest or
> a fully virtualized one

I certainly agree that we are not at a point where a final decision can be made. A major piece of that is that a layer 2 approach has not shown to be without a performance penalty.

A practical question. Do the IPs assigned to guests ever get used by anything besides the guest?

Eric

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
