Subject: Re: [RFC] network namespaces
Posted by Daniel Lezcano on Wed, 04 Oct 2006 09:40:32 GMT
View Forum Message <> Reply to Message

Andrey Savochkin wrote:
> Hi All,
>
> I'd like to resurrect our discussion about network namespaces.
> In our previous discussions it appeared that we have rather polar concepts
> which seemed hard to reconcile.
> Now I have an idea how to look at all discussed concepts to enable everyone's
> usage scenario.

Hi Andrey,

I have a few questions ... sorry for asking so late ;)

>
> 1. The most straightforward concept is complete separation of namespaces,
>   covering device list, routing tables, netfilter tables, socket hashes, and
>   everything else.
>
>   On input path, each packet is tagged with namespace right from the
>   place where it appears from a device, and is processed by each layer
>   in the context of this namespace.

If you have the namespace where is coming the packet, why do you tag the
packet instead of switching to the right namespace ?

>   Non-root namespaces communicate with the outside world in two ways: by
>   owning hardware devices, or receiving packets forwarded them by their parent
>   namespace via pass-through device.

Do you will do proxy arp and ip forwarding into the root namespace in
order to make non-root namespace visible from the outside world ?

Regards.

 -- Daniel