Subject: Re: Network namespaces a path to mergable code.
Posted by abdallah.chatila on Wed, 28 Jun 2006 05:59:18 GMT
View Forum Message <> Reply to Message

On Tue, Jun 27, 2006 at 10:33:48PM -0600, Eric W. Biederman wrote:
>
> Something to examine here is that if both network devices and sockets
> are tagged does that still allow implicit network namespace passing.

I think avoiding implicit network namespace passing expresses more
power/flexibility plus it would make things clearer to what
container/namespace a given network resource belongs too.

>From our experience with an implementation of network containers [Virtual
Routing for ipv4/ipv6, with a complete isolation between containers where ip
addresses can overlap...], there is some problem domain in which you cannot
afford to duplicate a process/daemon in each container [a big process for
instance, scalability w.r.t. number of containers etc]

By having a proper namespace tag per socket, this can be solved by allowing
a process running in the host context to create sockets in that namespace
than moving them to the target guest namespaces [via a special setsockopt
for instance or unix domain socket as you said].

Regards

>
> Eric
> -
> To unsubscribe from this list: send the line "unsubscribe netdev" in
> the body of a message to majordomo@vger.kernel.org
> More majordomo info at http://vger.kernel.org/majordomo-info.html