
Subject: Re: strict isolation of net interfaces

Posted by [Daniel Lezcano](#) on Fri, 30 Jun 2006 12:23:16 GMT

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

Serge E. Hallyn wrote:

> Quoting Cedric Le Goater (clg@fr.ibm.com):

>

>>we could work on virtualizing the net interfaces in the host, map them to
>>eth0 or something in the guest and let the guest handle upper network layers ?

>>

>>lo0 would just be exposed relying on skbuff tagging to discriminate traffic
>>between guests.

>

>

> This seems to me the preferable way. We create a full virtual net
> device for each new container, and fully virtualize the device
> namespace.

I have a few questions about all the network isolation stuff:

* What level of isolation is wanted for the network ? network devices
? IPv4/IPv6 ? TCP/UDP ?

* How is handled the incoming packets from the network ? I mean what
will be mechanism to dispatch the packet to the right virtual device ?

* How to handle the SO_BINDTODEVICE socket option ?

* Has the virtual device a different MAC address ? How to manage it
with the real MAC address on the system ? How to manage ARP, ICMP,
multicasting and IP ?

It seems for me, IMHO that will require a lot of translation and
browsing table. It will probably add a very significant overhead.

* How to handle NFS access mounted outside of the container ?

* How to handle ICMP_REDIRECT ?

Regards
