
Subject: Re: strict isolation of net interfaces

Posted by [Cedric Le Goater](#) on Fri, 30 Jun 2006 08:56:13 GMT

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

Serge E. Hallyn wrote:

>

> The last one in your diagram confuses me - why foo0:1? I would

> have thought it'd be

just thinking aloud. I thought that any kind/type of interface could be mapped from host to guest.

```
> host          | guest 0 | guest 1 | guest2
> -----+-----+-----+-----
> |          |          |          |
> |-> lo      <-----+--> lo0 ... | lo0      | lo0
> |          |          |          |
> |-> eth0     |          |          |
> |          |          |          |
> |-> veth0    <-----+--> eth0  |          |
> |          |          |          |
> |-> veth1    <-----+-----+-----+--> eth0
> |          |          |          |
> |-> veth2    <-----+-----+-----+--> eth0  |
>
```

> I think we should avoid using device aliases, as trying to do

> something like giving eth0:1 to guest1 and eth0:2 to guest2

> while hiding eth0:1 from guest2 requires some uglier code (as

> I recall) than working with full devices. In other words,

> if a namespace can see eth0, and eth0:2 exists, it should always

> see eth0:2.

>

> So conceptually using a full virtual net device per container

> certainly seems cleaner to me, and it seems like it should be

> simpler by way of statistics gathering etc, but are there actually

> any real gains? Or is the support for multiple IPs per device

> actually enough?

>

> Herbert, is this basically how ngnet is supposed to work?