
Subject: Re: strict isolation of net interfaces
Posted by [ebiederm](#) on Fri, 30 Jun 2006 17:41:59 GMT
[View Forum Message](#) <> [Reply to Message](#)

"Serge E. Hallyn" <serue@us.ibm.com> writes:

> Quoting Eric W. Biederman (ebiederm@xmission.com):
>> This whole debate on network devices show up in multiple network namespaces
>> is just silly. The only reason for wanting that appears to be better
> management.
>
> A damned good reason.

Better management is a good reason. But constructing the management in a way that hampers the implementation and confuses existing applications is a problem.

Things are much easier if namespaces are completely independent.

Among other things the semantics are clear and obvious.

> Clearly we want the parent namespace to be able
> to control what the child can do. So whatever interface a child gets,
> the parent should be able to somehow address. Simple iptables rules
> controlling traffic between it's own netdevice and the one it hands it's
> children seem a good option.

That or we setup the child and then drop CAP_NET_ADMIN.

>> We have deeper issues like can we do a reasonable implementation without a
>> network device showing up in multiple namespaces.
>
> Isn't that the same issue?

I guess I was thinking from the performance and cleanliness point of view.

>> If we can get layer 2 level isolation working without measurable overhead
>> with one namespace per device it may be worth revisiting things. Until
>> then it is a side issue at best.
>
> Ok, and in the meantime we can all use the network part of the bsdjail
> lsm? :)

If necessary. But mostly we concentrate on the fundamentals and figure out what it takes to take the level 2 stuff working.

Eric
