Subject: Re: Network virtualization/isolation
Posted by ebiederm on Tue, 28 Nov 2006 16:51:57 GMT
View Forum Message <> Reply to Message

I do not want to get into a big debate on the merits of various techniques at this time. We seem to be in basic agreement about what we are talking about.

There is one thing I think we can all agree upon.
- Everything except isolation at the network device/L2 layer, does not allow guests to have the full power of the linux networking stack.

- There has been a demonstrated use for the full power of the linux networking stack in containers..

- There are a set of techniques which look as though they will give us full speed when we do isolation of the network stack at the network device/L2 layer.

Is there any reason why we don't want to implement network namespaces without the full power of the linux network stack?

If there is a case where we clearly don't want the full power of the linux network stack in a guest but we still need a namespace we can start looking at the merits of the alternatives.

> What is this new paradigm you are talking about ?

The basic point is this. The less like stock linux the inside of a container looks, and the more of a special case it is the more confusing it is. The classic example is that for a system container routing packets between containers over the loopback interface is completely unexpected.

> There is not extra networking data structure instantiation in the
> Daniel's L3.
Nope just an extra field which serves the same purpose.

>> - Bind/Connect/Accept filtering. There are so few places in
>> the code this is easy to maintain without sharing code with
>> everyone else.
>
> For isolation too ? Can we build network migration on top of that ?

As long as you can take your globally visible network address with you when you migrate you can build network migration on top of it. So yes bind/accept filtering is sufficient to implement migration, if you are

only using IP based protocols.

Eric

_____

Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers