Subject: Re: [RFD] L2 Network namespace infrastructure
Posted by ebiederm on Sun, 24 Jun 2007 12:38:38 GMT
View Forum Message <> Reply to Message

David Miller <davem@davemloft.net> writes:

> From: ebiederm@xmission.com (Eric W. Biederman)
> Date: Sat, 23 Jun 2007 16:56:49 -0600
>
>> If the only use was strong isolation which Dave complains about I would
>> concur that the namespace approach is inappropriate. However there are
>> a lot other uses.
>
> By your very admission the only appropriate use case is when users
> are not "hostile" and can be trusted to some extent.

Yes. Like all of Linux. Totally hostile antagonistic users that
have millions or billions of dollars to spend better figuring
out how to mess each other up and do bad things to each other should
not be running code on the same machine.

> And that by definition makes it not appropriate for a general purpose
> operating system like Linux.

Not at all. The security should be at least as good as between different
users today, and most likely better.

> Containers are I believe a step backwards, and we're better than that.

I heartily disagree.

I do agree that if someone can write a minimal hypervisor it is possible
to provide stronger guarantees of separation on the same hardware then
it is with linux. But that is because a hypervisor should be such a small
code base one man should be able to prove and audit the thing. Which allows
several independent people to reproduce that effort.

So my real admission is that hypervisor should be able to do much
stronger isolation.

Namespaces come in when you want to do silly little things like share the
sysadmin responsibility, or not modify applications that make silly assumptions
or don't want to take the huge resource consumption hit hypervisors require.

Eric

_____
Containers mailing list
Containers@lists.linux-foundation.org