
Subject: Re: [RFC] network namespaces

Posted by [dev](#) on Wed, 06 Sep 2006 14:52:50 GMT

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

>>On Tue, Sep 05, 2006 at 08:45:39AM -0600, Eric W. Biederman wrote:

>>

>>>Daniel Lezcano <dlezcano@fr.ibm.com> writes:

>>>

>>>For HPC if you are interested in migration you need a separate IP

>>>per container. If you can take you IP address with you migration of

>>>networking state is simple. If you can't take your IP address with you

>>>a network container is nearly pointless from a migration perspective.

>>>

>>>Beyond that from everything I have seen layer 2 is just much cleaner

>>>than any layer 3 approach short of Serge's bind filtering.

>>

>>well, the 'ip subset' approach Linux-VServer and

>>other Jail solutions use is very clean, it just does

>>not match your expectations of a virtual interface

>>(as there is none) and it does not cope well with

>>all kinds of per context 'requirements', which IMHO

>>do not really exist on the application layer (only

>>on the whole system layer)

>

>

> I probably expressed that wrong. There are currently three

> basic approaches under discussion.

> Layer 3 (Basically bind filtering) nothing at the packet level.

> The approach taken by Serge's version of bsdjails and Vserver.

>

> Layer 2.5 What Daniel proposed.

>

> Layer 2. (Trivially mapping each packet to a different interface)

> And then treating everything as multiple instances of the

> network stack.

> Roughly what OpenVZ and I have implemented.

I think classifying network virtualization by Layer X is not good enough.

OpenVZ has Layer 3 (venet) and Layer 2 (veth) implementations, but

in both cases networking stack inside VE remains fully virtualized.

Thanks,

Kirill

Containers mailing list

Containers@lists.osdl.org

<https://lists.osdl.org/mailman/listinfo/containers>
