
Subject: Re: [RFC] network namespaces
Posted by [ebiederm](#) on Wed, 16 Aug 2006 17:35:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

Alexey Kuznetsov <kuznet@ms2.inr.ac.ru> writes:

> Hello!
>
>> (application) containers. Performance aside, are there any reasons why
>> this approach would be problematic for c/r?
>
> This approach is just perfect for c/r.

Yes. For c/r you need to take your state with you.

> Probably, this is the only approach when migration can be done
> in a clean and self-consistent way.

Basically there are currently 3 approaches that have been proposed.

The trivial bsdjail style as implemented by Serge and in a slightly more sophisticated version in vserver. This approach as it does not touch the packets has little to no packet level overhead. Basically this is what I have called the Level 3 approach.

The more in depth approach where we modify the packet processing based upon which network interface the packet comes in on, and it looks like each namespace has it's own instance of the network stack. Roughly what was proposed earlier in this thread the Level 2 approach. This potentially has per packet overhead so we need to watch the implementation very carefully.

Some weird hybrid as proposed by Daniel, that I was never clear on the semantics.

>From the previous conversations my impression was that as long as we could get a Layer 2 approach that did not slow down the networking stack and was clean, everyone would be happy.

I'm buried in the process id namespace at the moment, and except to be so for the rest of the month, so I'm not going to be very helpful except for a few stray comments.

Eric
