
Subject: RE: [RFC] network namespaces

Posted by [Caitlin Bestler](#) on Wed, 06 Sep 2006 23:06:16 GMT

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

ebiederm@xmission.com wrote:

>
>> Finally, as I understand both network isolation and network
>> virtualization (both level2 and level3) can happily co-exist. We do
>> have several filesystems in kernel. Let's have several network
>> virtualization approaches, and let a user choose. Is that makes
>> sense?
>
> If there are not compelling arguments for using both ways of
> doing it is silly to merge both, as it is more maintenance overhead.
>

My reading is that full virtualization (Xen, etc.) calls for
implementing
L2 switching between the partitions and the physical NIC(s).

The tradeoffs between L2 and L3 switching are indeed complex, but
there are two implications of doing L2 switching between partitions:

- 1) Do we really want to ask device drivers to support L2 switching for
partitions and something *different* for containers?
- 2) Do we really want any single packet to traverse an L2 switch (for
the partition-style virtualization layer) and then an L3 switch
(for the container-style layer)?

The full virtualization solution calls for virtual NICs with distinct
MAC addresses. Is there any reason why this same solution cannot work
for containers (just creating more than one VNIC for the partition,
and then assigning each VNIC to a container?)
