
Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by [Ben Greear](#) on Mon, 26 Jun 2006 22:13:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:

> Basically it is just a matter of:
> if (dest_mac == my_mac1) it is for device 1.
> If (dest_mac == my_mac2) it is for device 2.
> etc.
>
> At a small count of macs it is trivial to understand it will go
> fast for a larger count of macs it only works with a good data
> structure. We don't hit any extra cache lines of the packet,
> and the above test can be collapsed with other routing lookup tests.

I think you should do this at the layer-2 level, well before you get to routing. That will make the virtual mac-vlan work with arbitrary protocols and appear very much like a regular ethernet interface. This approach worked well with .1q vlans, and with my version of the mac-vlan module.

Using the mac-vlan and source-based routing tables, I can give a unique 'interface' to each process and have each process able to bind to the same IP port, for instance. Using source-based routing (by binding to a local IP explicitly and adding a route table for that source IP), I can give unique default routes to each interface as well. Since we cannot have more than 256 routing tables, this approach is currently limited to around 250 virtual interfaces, but that is still a substantial amount.

My mac-vlan patch, redirect-device patch, and other hackings are consolidated in this patch:

http://www.candelatech.com/oss/candela_2.6.16.patch

Thanks,
Ben

--
Ben Greear <greearb@candelatech.com>
Candela Technologies Inc <http://www.candelatech.com>
