
Subject: Re: Network virtualization/isolation
Posted by [jamal](#) on Sun, 03 Dec 2006 16:58:30 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Sun, 2006-03-12 at 17:37 +0100, Herbert Poetzl wrote:
> On Sun, Dec 03, 2006 at 07:26:02AM -0500, jamal wrote:

> To use an extreme example: if i picked apache as a
> > binary compiled 10 years ago, it will run on the L2 approach but not on
> > the L3 approach. Is this understanding correct? I find it hard to
> > believe that the L3 approach wouldnt work this way - it may be just my
> > reading into the doc.
>
> the 10 year old apache will run with layer 3 isolation
> as well as with layer 2 virtualization (probably a little
> faster though, we do not know yet :), because what it
> does is IP (layer 3) traffic ...
>

Ok, thanks for clarifying this.

> > I think the litmus test for this approach is the answer to the question:
> > If i compiled in the containers in and do not use the namespaces, how
> > much more overhead is there for the host path? I would hope that it is
> > as close to 0 as possible. It should certainly be 0 if i dont compile in
> > containers.
>
> IMHO there are three cases to consider, to get valid
> 'performance' numbers:
>
> - host system with and without containers enabled
> - single guest (container) compared to host system _without_

Sound reasonable.

> - bunch of guests (e.g. 10) compared to 10 apps/threads on host
>

Your mileage may vary. For me trying to run virtual routers; this is not an important test. I want to be able to have containers each running quagga and OSPF. I cant achieve my goals with with 10 quaggas without making some major changes to quagga.

> one proven feature of the L3 isolation is that those
> all end up with the same or even better performance

I think it is valuable to reduce the overhead. I think that it may be reasonable to some threshold to trade a little performance for

genericity. What the threshold is, i dont know.

- > > - Manageability from the host side. It seems to be more complex with the
- > > L2 than with L3. But so what? These tools are written from scratch and
- > > there is no "backward compatibility" baggage.
- >
- > well, no, actually the 'tools' to manage layer 3 isolation
- > are already there,
- > and except for the 'setup' there is
- > nothing special to configure, as networking still lives
- > on the host
- >

I dont see the two as being separate issues. You must create container; you must configure networking on them; it is forgivable to have the second part of that process to involve some non-standard tools for the containers (from the host).

It is not forgivable to have speacilized tools within the container.

- > I would be interested in a config layout for a typical
- > L3 isolation setup when you 'only' have L2 virtualization
- >
- > - typical host system with apache, mysql, postfix, ssh
- > and ftp is broken down into security contexts to
- > allow for increased security
- > - as part of that process, the services are isolated,
- > while apache and ftp share the same ip [ip0], mysql
- > will be using a local one [ip1], and postfix/ssh a
- > second public one [ip2]
- >
- > the L3 isolation approach is straight forward:
- >
- > - assign the two public ips to eth0, the local one
- > to lo or dummy0
- > - create five isolation areas where 0 and 1 share ip0,
- > 2 uses ip1 and 3,4 uses ip2
- >
- > that's it, all will work as expected ... let's see with
- > what L2 isolation example you come up with, which is
- > able to 'mimic' this setup ...
- >
- > note: no question it is possible to do that with L2

Unless i am misreading, isnt this merely a matter of configuring on the container side eth0 (I think you are talking about VE side eth0 in your example above) two public ip addresses (or even two ethx devices) and then attach IP addresses to them? mysql gets an lo address. Would this not work?

Out of curiosity: assume we have a local LAN (perhaps something upstream does NAT), is it possible to have the same IP address going to multiple containers?

cheers,
jamal

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
