

---

Subject: Re: L2 network namespace benchmarking (resend with Service Demand)  
Posted by [Benjamin Thery](#) on Fri, 06 Apr 2007 11:19:36 GMT

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

---

Eric W. Biederman wrote:

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

>

>> Hi,

>>

>> as suggested Rick, I added the Service Demand results to the matrix.

>

> A couple of random thoughts in trying to understand the numbers you are  
> seeing.

>

> - Checksum offloading?

>

> You have noted that with the bridge netfilter support disabled you  
> are still seeing additional checksum overhead. Just like you are  
> seeing in the routing case.

>

> Is it possible the problem is simply that etun doesn't support  
> checksum offloading, while your normal test hardware does?

Looks like you are 100% correct.

I feel a bit stupid I didn't think about this "small" difference  
between real NIC and etun.

If I turn off checksum offloading on my physical NIC, the checksum  
"overhead" (load) measured by oprofile is about the same in both case:  
when running netperf through a real NIC or through an etun tunnel first.

Benjamin

> - Tagged VLANs?

>

> Currently you have tested bridging and routing to get the packets to  
> a network namespace. Could you test tagged vlans?

>

> I'm just curious if we have anything in the network stack today that  
> will multiplex a NIC without measurable overhead.

>

> - Without NETNS?

>

> We should probably see if we can setup the same configuration we are  
> testing without network namespaces (just multiple interfaces on the  
> same machine) and see if we can still measure the same overhead.  
> Just to confirm the overhead is not a network namespace related  
> thing.

>  
> I know we can configure the same case with bridging and I am fairly  
> confident that we will see the same overhead without network  
> namespaces.  
>  
> Of the top of my head I am insufficiently clever to think how we  
> could configure the routing case without network namespaces,  
> although we might be able to force it and if so it would be  
> interesting to measure.  
>  
> I will work to get the etun setup races fixed and to fix whatever  
> obvious feature deficiencies it has (like no configurable MTU support)  
> and see if I can get that pushed upstream. That should make it easier  
> for other people to reproduce what we are seeing.  
>  
> Eric  
>  
> \_\_\_\_\_  
> Containers mailing list  
> Containers@lists.linux-foundation.org  
> <https://lists.linux-foundation.org/mailman/listinfo/containers>  
>

--  
Benjamin Thery - BULL/DT/Open Software R&D

<http://www.bull.com>

\_\_\_\_\_  
Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---