

Benjamin Thery <benjamin.thery@bull.net> writes:

> Eric W. Biederman wrote:

>> 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.

Interesting. You can also 'enable' checksum offloading when using etun with
ethtool. Which should just tell the kernel not to do checksumming. A
bad idea in general but it might be useful in confirming where the
performance overhead is coming from, and when used with routing I
believe it is safe. When used with bridging I don't know.

Thinking about it the ideal situation is to preserve `skb->ip_summed` if it
came from another device, instead of unconditionally setting it.
I need to take a good hard look at `etun_xmit` and make certain we
are dotting all of the i's and crossing all of the t's for best
performance and compatibility with the rest of the network stack.

Eric

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