Subject: Re: Re: [PATCHSET] 2.6.20-lxc8
Posted by Benjamin Thery on Wed, 28 Mar 2007 12:30:26 GMT
View Forum Message <> Reply to Message

Eric W. Biederman wrote:
> Kirill Korotaev <dev@sw.ru> writes:
>
>> Benjamin,
>>
>> checksumming can be optimized out as well.
>> We had an experimental patch for OpenVZ venet device, which adds
>> NETIF_F_LLTX | NETIF_F_HW_CSUM | NETIF_F_SG | NETIF_F_HIGHDMA
>> features to venet device and avoids additional checksumming where possible
>> (moving RX/TX checksum calculation to hardware).
>>
>> So I guess this is doable in future as well.
>
> I think I have the checksum bits settable in software with etun already. If not
> it shouldn't be to hard to add.

I tried to activate the checksum offload on etun to see if it improves
things, but, unfortunately once activated all my traffic was lost or
blocked. I didn't spend a lot of time on the issue. May be I'll give
it another try.

BTW, there is a small bug in etun_set_tx_csum. We can't disable the
checksum offloading once it has been set. I think it should look like
this: (sorry I haven't a patch ready)

```
static int etun_set_tx_csum(struct net_device *dev, u32 data)
{
 if (data)
  dev->features |= NETIF_F_NO_CSUM;
 else
  dev->features &= ~NETIF_F_NO_CSUM;
 return 0;
}
```

>
> I don't default to that because depending on your configuration it might not
> be safe. In particular I think when you are using ethernet bridging we
> need to do the packet checksum immediately off the wire.
>
> Eric
>

--
B e n j a m i n T h e r y - 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