Subject: Re: [PATCH] net: Add etun driver
Posted by Roland Dreier on Fri, 06 Apr 2007 20:57:25 GMT
View Forum Message <> Reply to Message

```
> +/*
 > + * The higher levels take care of making this non-reentrant (it's
 > + * called with bh's disabled).
 > + */
 > +static int etun_xmit(struct sk_buff *skb, struct net_device *tx_dev)
```

You have this comment, but then...

```
 > + dev->features  = NETIF_F_FRAGLIST
 > +     | NETIF_F_HIGHDMA
 > +     | NETIF_F_LLTX;
```

you set LLTX, which means that the upper layers _don't_ make sure that
your xmit routine is not reentrant.

It looks like the impact of multiple simultaneous xmit calls is just
the possibility of screwing up the statistics, but still I think you
want to drop the LLTX feature (since you have no lock of your own to
try and take).

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

Subject: Re: [PATCH] net: Add etun driver
Posted by ebiederm on Sat, 07 Apr 2007 02:08:22 GMT
View Forum Message <> Reply to Message

Roland Dreier <rdreier@cisco.com> writes:

```
> > +/*
> > + * The higher levels take care of making this non-reentrant (it's
> > + * called with bh's disabled).
> > + */
> > +static int etun_xmit(struct sk_buff *skb, struct net_device *tx_dev)
>
> You have this comment, but then...
>
> > + dev->features  = NETIF_F_FRAGLIST
> > +     | NETIF_F_HIGHDMA
> > +     | NETIF_F_LLTX;
>
> you set LLTX, which means that the upper layers _don't_ make sure that
```

> your xmit routine is not reentrant.
>
> It looks like the impact of multiple simultaneous xmit calls is just
> the possibility of screwing up the statistics, but still I think you
> want to drop the LLTX feature (since you have no lock of your own to
> try and take).

Yup. That is an inconsistency and probably a bug. I have to think
through what makes most sense in this case. Though the cheap answer
is clearly to remove NETIF_F_LLTX.