

Hi !

Eric W. Biederman wrote:

[ ... ]

> So just to sink one additional nail in the coffin of the silly  
> guest to guest communication issue. For any two guests where  
> fast communication between them is really important I can run  
> an additional interface pair that requires no routing or bridging.  
> Given that the implementation of the tunnel device is essentially  
> the same as the loopback interface and that I make only one  
> trip through the network stack there will be no performance overhead.  
> Similarly for any critical guest communication to the outside world  
> I can give the guest a real network adapter.  
>  
> That said I don't think those things will be necessary and that if  
> they are it is an optimization opportunity to make various bits  
> of the network stack faster.

just one comment on the 'guest to guest communication' topic :

guest to guest communication is an important factor in consolidation scenarios, where containers are packed on one server. This for maintenance issues or priority issues on a HPC cluster for example. This case of container migration is probably the most interesting and the performance should be more than acceptable. May be not a top priority for the moment.

thanks,

C.

---