Subject: Re: Re: network namespace ipv6 perfs
Posted by Benjamin Thery on Mon, 03 Mar 2008 15:04:47 GMT
View Forum Message <> Reply to Message

On Mon, Mar 3, 2008 at 3:55 PM, Pavel Emelyanov <xemul@openvz.org> wrote:
> Benjamin Thery wrote:
> > Daniel Lezcano wrote:
> >> Hi,
> >>
> >> Some performance tests was made by Benjamin to watch out the impact of
> >> the network namespace. The good news is there is no impact when used
> >> with or without namespaces. That has been checked using a real network
> >> device inside a network namespace.
> >>
> >> These results are consistent with the ones previously made for ipv4.
> >>
> >> http://lxc.sourceforge.net/network/bench_ipv6_graph.php
> >>
> >> Thanks to Benjamin who did all the performance tests :)
> >
> > In these results, may be, there is one thing that should be explained.
> > It is the CPU utilization overhead in the 'veth' case.
> >
> > Compared to physical devices or macvlan, veth interfaces don't benefit
> > from hardware offloading mechanisms: i.e. checksums have to be computed
> > by the soft. That explains the big overhead in CPU utilization when
>
> You can tune the veth devices not to account checksum when unnecessary.

Oh. This is interesting.

You mean with ethtool -K rx/tx?
I will give it a try.

Benjamin

>
>
> > using this kind of virtual interface.
> >
> > Benjamin
> >
> >> Regards
> >>   -- Daniel
> >>

>--
> To unsubscribe from this list: send the line "unsubscribe netdev" in

> the body of a message to majordomo@vger.kernel.org
> More majordomo info at http://vger.kernel.org/majordomo-info.html
>

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