
Subject: Re: network namespace ipv6 perfs
Posted by [Daniel Lezcano](#) on Mon, 03 Mar 2008 20:01:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

Rick Jones 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.

>

> The *_RR tests seem to show a drop in throughput and corresponding
> increases in service demand - could that be because things like TSO et
> al cannot mask much of anything in the way of a path-length increase?

Hmm. In fact Benjamin took the 2.6.23.16 kernel where there were no network namespace code at all. So these differences between 2.6.23.16 and 2.6.25-rc1 does not show a performance degradation especially related to the network namespaces. The important point is the 2.6.25-rc1 without ipv6 netns and 2.6.25-rc1 with ipv6 netns code applied, I mean the second and the third line and we can point that the ipv6 netns code does not degrade performances for either throughput and service demand.

> From the annotations, I'm ass-u-me-ing that NS was only used on the
> netperf side and not both netperf and netserver side?

right :)

> happy benchmarking,

Thanks Rick.

-- Daniel

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