Subject: Re: L2 network namespace benchmarking
Posted by Rick Jones on Wed, 28 Mar 2007 18:08:28 GMT
View Forum Message <> Reply to Message

> If I read the results right it took a 32bit machine from AMD with
> a gigabit interface before you could measure a throughput difference.
> That isn't shabby for a non-optimized code path.

Just some paranoid ramblings - one needs to look beyond just whether or
not the performance of a bulk transfer test (eg TCP_STREAM) remains able
to hit link-rate. One has to also consider the change in service demand
(the normalization of CPU util and throughput). Also, with
functionality like TSO in place, the ability to pass very large things
down the stack can help cover for a multitude of path-length sins. And
with either multiple 1G or 10G NICs becoming more and more prevalent, we
have another one of those "NIC speed vs CPU speed" switch-overs, so
maintaining single-NIC 1 gigabit throughput, while necessary, isn't
(IMO) sufficient.

Soooo, it becomes very important to go beyond just TCP_STREAM tests when
evaluating these sorts of things. Another test to run would be the
TCP_RR test. TCP_RR with single-byte request/response sizes will
"bypass" the TSO stuff, and the transaction rate will be more directly
affected by the change in path length than a TCP_STREAM test. It will
also show-up quite clearly in the service demand. Now, with NICs doing
interrupt coalescing, if the NIC is strapped "poorly" (IMO) then you may
not see a change in transaction rate - it may be getting limited
artifically by the NIC's interrupt coalescing. So, one has to fall-back
on service demand, or better yet, disable the interrupt coalescing.

Otherwise, measuring peak aggregate request/response becomes necessary.

rick jones
don't be blinded by bit-rate
_____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers