
Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view
Posted by [Ben Greear](#) on Tue, 27 Jun 2006 16:07:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

Ben Greear wrote:

> Herbert Poetzl wrote:

>

>> On Mon, Jun 26, 2006 at 03:13:17PM -0700, Ben Greear wrote:

>

>

>> yes, that sounds good to me, any numbers how that

>> affects networking in general (performance wise and

>> memory wise, i.e. caches and hashes) ...

>

>

> I'll run some tests later today. Based on my previous tests,

> I don't remember any significant overhead.

Here's a quick benchmark using my redirect devices (RDD). Each RDD comes in a pair...when you tx on one, the pkt is rx'd on the peer. The idea is that it is exactly like two physical ethernet interfaces connected by a cross-over cable.

My test system is a 64-bit dual-core Intel system, 3.013 Ghz processor with 1GB RAM. Fairly standard stuff..it's one of the Shuttle XPC systems. Kernel is 2.6.16.16 (64-bit).

Test setup is: rdd1 -- rdd2 [bridge] rdd3 -- rdd4

I am using my proprietary module for the bridge logic...and the default bridge should be at least this fast. I am injecting 1514 byte packets on rdd1 and rdd4 with pktgen (bi-directional flow). My pktgen is also receiving the pkts and gathering stats.

This setup sustains 1.7Gbps of generated and received traffic between rdd1 and rdd4.

Running only the [bridge] between two 10/100/1000 ports on an Intel PCI-E NIC will sustain about 870Mbps (bi-directional) on this system, so the virtual devices are quite efficient, as suspected.

I have not yet had time to benchmark the mac-vlans...hopefully later today.

Thanks,
Ben

--

Ben Greear <greearb@candelatech.com>
Candela Technologies Inc <http://www.candelatech.com>
