

Re: em forwarding performance (was Proposed 6.2 em RELEASE patch)

Re: em forwarding performance (was Proposed 6.2 em RELEASE patch)

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2006-11/msg00065.html>

- *From:* Mike Tancsa <mike@xxxxxxxxxx>
 - *Date:* Thu, 30 Nov 2006 13:04:08 -0500
-

At 12:57 PM 11/30/2006, Ivan Voras wrote:

Mike Tancsa wrote:

- > Yeah I inadvertently slighted the NetBSD folks by leaving them out. So
- > I guess I better give them a try as well.
- >
- > The part that really surprises me is the drop in performance as firewall
- > rules are added to RELENG_6 and above. Both LINUX and RELENG_4 seem to
- > scale well with the number of rules added but RELENG_6 takes a big drop.

Wasn't there some important setting in ipfw you can tweak if you need lots of ipfw rules? Size of some hash table?

Quick Googling found this: http://info.iet.unipi.it/~luigi/ip_dummynet/ and net.inet.ip.fw.dyn_buckets: 256. AFAIK the hash size needed to be tweaked manually in the code, and net.inet.ip.fw.dyn_buckets: 256 is listed as read-only so this might be it. Maybe mailing Luigi will help finding out...

I was told offlist "there is additional per-packet locking overhead not seen in RELENG_4 where all processing is covered by the same spl."...

---Mike

freebsd-performance@xxxxxxxxxxxxx mailing list
<http://lists.freebsd.org/mailman/listinfo/freebsd-performance>

To unsubscribe, send any mail to "freebsd-performance-unsubscribe@xxxxxxxxxxxxx"