

Re: Avoiding natd overhead

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2006-10/msg00214.html>

- *From:* Julian Elischer <julian@xxxxxxxxxxxxx>
 - *Date:* Sat, 21 Oct 2006 21:21:15 -0700
-

Brett Glass wrote:

At 09:50 PM 10/21/2006, Julian Elischer wrote:

one thing that you need to name sure of is that only the packets that have potential of being on interest to natd are passed to natd.

I do. In fact, this is a capability I would lose if I used ipfilters or pf to do NAT, which is why I want to find a way to use a mechanism that's triggered by IPFW.

You were the person who invented "divert sockets," were you not? How hard would it be to create a mechanism (a sort of "kernel divert socket") so that kernel modules and/or netgraph nodes could do the same things which are now done by userland processes listening on divert sockets? This would boost the performance of any FreeBSD machine that did NAT (which many if not most do).

you can in two ways..

create a netgraph ksocket node of type divert
then attach that to a netgraph ng_nat node.

OR in 7.0 you can call netgraph directly

there is a netgraph keyword in ipfw.

--Brett Glass

freebsd-net@xxxxxxxxxxxxx mailing list
<http://lists.freebsd.org/mailman/listinfo/freebsd-net>
To unsubscribe, send any mail to "freebsd-net-unsubscribe@xxxxxxxxxxxxx"