

freebsd-net: Re: Newbie question: is net.inet.ip.fastforwarding intended for routersonly?

Re: Newbie question: is net.inet.ip.fastforwarding intended for routersonly?

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2004-01/0196.html>

From: Andre Oppermann (*andre_at_freebsd.org*)

Date: 01/17/04

Date: Sat, 17 Jan 2004 18:53:46 +0100

To: Xin LI <delphij@frontfree.net>

Xin LI wrote:

>

> *As title. To my imagination (I have taken a sight on the kernel networking
> code), the fastforwarding path is intended for routers only, so if I want
> some functionalities, for example, NAT, the fastforwarding is not useful,
> and my experiment shows that if I enable it, ipfilter's NAT will be broken.
> Am I right?*

I have written a new fastforward code for 5.2 and -current which does everything including ipfw, NAT and ipfilter.

Here is the blurb for it:

/*

* ip_fastforward gets its speed from processing the forwarded packet to
* completion (if_output on the other side) without any queues or netisr's.
* The receiving interface DMA's the packet into memory, the upper half of
* driver calls ip_fastforward, we do our routing table lookup and directly
* send it off to the outgoing interface which DMA's the packet to the
* network card. The only part of the packet we touch with the CPU is the
* IP header (unless there are complex firewall rules touching other parts
* of the packet, but that is up to you). We are essentially limited by bus
* bandwidth and how fast the network card/driver can set up receives and
* transmits.

*

* We handle basic errors, ip header errors, checksum errors,
* destination unreachable, fragmentation and fragmentation needed and
* report them via icmp to the sender.

*

* Else if something is not pure IPv4 unicast forwarding we fall back to
* the normal ip_input processing path. We should only be called from
* interfaces connected to the outside world.

*

* Firewalling is fully supported including divert, ipfw fwd and ipfilter
* ipnat and address rewrite.

Re: Newbie question: is net.inet.ip.fastforwarding intended for routersonly?

freebsd-net: Re: Newbie question: is net.inet.ip.fastforwarding intended for routersonly?

*

* IPSEC is not supported if this host is a tunnel broker. IPSEC is supported for connections to/from local host.

*

* We try to do the least expensive (in CPU ops) checks and operations first to catch junk with as little overhead as possible.

*

* We take full advantage of hardware support for ip checksum and fragmentation offloading.

*

* We don't do ICMP redirect in the fast forwarding path. I have had my own cases where two core routers with Zebra routing suite would send millions of ICMP redirects to connected hosts if the router to dest was not the default gateway. In one case it was filling the routing table of a host with close to 300'000 cloned redirect entries until it ran out of kernel memory. However the networking code proved very robust and it didn't crash or went ill otherwise.

*/

--

Andre

freebsd-net@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-net>

To unsubscribe, send any mail to "freebsd-net-unsubscribe@freebsd.org"