

Re: tcpdump – tun/tap virtual interfaces

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2003-09/0171.html>

From: Robert Watson (rwatson_at_freebsd.org)

Date: 09/28/03

Date: Sun, 28 Sep 2003 14:05:24 -0400 (EDT)

To: "Giovanni P. Tirloni" <gpt@tirloni.org>

On Sun, 28 Sep 2003, Giovanni P. Tirloni wrote:

> * Robert Watson (rwatson@freebsd.org) wrote:
>
>> *Do you see anything when you ping the broadcast address or other foreign
>> address of the tap interface? Packets delivered to local IP addresses
>> generally don't go out an interface.*
>
> *About Ethernet frames not going out to the wire and being sent to the
> loopback..*
>
> *The check seems to happen at line 291 in if_ethersubr.c and then it
> uses the if_simloop() function to copy the packet to the loopback
> interface. Is that right?*
>
> *The rcvif interface is set to the hardware device, how is this used in
> this case? What kind of checks are done to the rcvif usually?*
>
> *I haven't received my copy of Steven's Volume 2 yet so if it's
> explained there (as I hope) I will sit in my corner and wait to for it
> patiently :)*

Ethernet loopback does occur, and BPF will pick those up. However, the loopback you're seeing is actually happening at the IP layer, as a result of routing rather than link layer behavior:

```
10 link#6 UC 1 0 tap0
10.0.10.1 00:bd:18:a1:11:00 UHLW 0 26 lo0
```

Local IP addresses have their packets routed to them over lo0, so the packets being looked for can be found by doing tcpdump on lo0:

```
test1# tcpdump -eni lo0 &
[2] 511
tcpdump: listening on lo0
test1# Sep 28 14:03:07 test1 kernel: lo0: promiscuous mode enabled
```

freebsd-net: Re: tcpdump – tun/tap virtual interfaces

```
test1# ping -c 1 10.0.10.1
PING 10.0.10.1 (10.0.10.1): 56 data bytes
64 bytes from 10.0.10.1: icmp_seq=0 ttl=64 time=0.073 ms

--- 10.0.10.1 ping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.073/0.073/0.073/0.000 ms
test1# 14:03:12.713690 AF 2 84: 10.0.10.1 > 10.0.10.1: icmp: echo request
14:03:12.713724 AF 2 84: 10.0.10.1 > 10.0.10.1: icmp: echo reply
```

Route command output appended below.

Robert N M Watson FreeBSD Core Team, TrustedBSD Projects
robert@fledge.watson.org Network Associates Laboratories

```
route get 10.0.10.1
  route to: 10.0.10.1
destination: 10.0.10.1
interface: lo0
  flags: <UP,HOST,DONE,LLINFO,WASCLONED,LOCAL>
rcvpipe sndpipe ssthresh rtt,msec rttvar hopcount mtu
expire
  0 0 0 0 0 0 1500
0
test1# route get 10.0.10.2
  route to: 10.0.10.2
destination: 10.0.0.0
  mask: 255.0.0.0
interface: tap0
  flags: <UP,DONE,CLONING>
rcvpipe sndpipe ssthresh rtt,msec rttvar hopcount mtu
expire
  0 0 0 0 0 0 1500
-100
```

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"