

## DragonFlyBSD - Bug #3045

### net.inet.ip.fastforwarding causes PF NAT to break/duplicate packets

06/18/2017 04:26 PM - benjolitz

|                        |               |                        |            |
|------------------------|---------------|------------------------|------------|
| <b>Status:</b>         | Resolved      | <b>Start date:</b>     | 06/18/2017 |
| <b>Priority:</b>       | Normal        | <b>Due date:</b>       |            |
| <b>Assignee:</b>       |               | <b>% Done:</b>         | 0%         |
| <b>Category:</b>       | Networking    | <b>Estimated time:</b> | 0.00 hour  |
| <b>Target version:</b> | Latest stable |                        |            |

#### Description

This issue took two days to track down.

Here is the gist of the issue:

If you set net.inet.ip.fastforwarding = 1  
AND set your pf.conf to:

```
int_if="int1"  
ext_if="ext0"  
localnet=$int_if:network  
nat pass on $ext_if inet from $localnet to any -> ($ext_if)
```

Your tcpdump (for a telnet google.com 80 with `GET /` and a new line entered into the buffer) will show retransmits and unbelievably slow NAT behavior (we're talking about one blob of 120 bytes every 30 to 60 seconds)

This bug is filed in hopes that google will pick up the following terms:  
pf nat slow, pf nat tcp duplicate

If fastforwarding is disabled, PF will behave correctly.

Here is a sample of it:

```
14:17:46.365163 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 3933726782, win 4104, options [nop,nop,TS val 766934808 ecr 932753216], length 0  
E..4?.@.?..2.....W.P.f%z.w.>.....  
-...7..@  
14:17:47.171679 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 1, win 4104, options [nop,nop,TS val 766935613 ecr 932753216], length 0  
E..4vE@.?..P.....W.P.f%z.w.>.....  
-..=7..@  
14:17:47.412131 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.], seq 0:7, ack 1, win 4104, options [nop,nop,TS val 766935853 ecr 932753216], length 7  
E.;.@.?.....W.P.f%z.w.>....(.....  
-..-7..@GET /  
14:17:47.556687 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.], seq 0:7, ack 1, win 4104, options [nop,nop,TS val 766935998 ecr 932753216], length 7  
E.;;Z@.?.....W.P.f%z.w.>....'.....  
-...7..@GET /  
14:17:47.708820 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.], seq 0:7, ack 1, win 4104, options [nop,nop,TS val 766936150 ecr 932753216], length 7  
E.;;9@.?.....W.P.f%z.w.>....'_.....  
-..V7..@GET /  
14:17:48.104957 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.], seq 0:9, ack 1, win 4104, options [nop,nop,TS val 766936542 ecr 932753216], length 9  
E..=..@.?..A  
.....W.P.f%z.w.>.....  
-...7..@GET /  
14:17:48.698278 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.], seq 0:9, ack 1, win 4104, options
```

```
[nop,nop,TS val 766937127 ecr 932753216], length 9
E..=H@.?.....:W.P.f%z.w.>.....
-..7..@GET /

14:17:48.700218 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 1, win 4104, options [nop,nop,TS val
766937129 ecr 932753216], length 0
E..4..@.?.....:W.P.f%.w.>.....
-..7..@
14:17:49.668670 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.], seq 0:9, ack 1, win 4104, options
[nop,nop,TS val 766938095 ecr 932753216], length 9
E..=..@.?..c.....:W.P.f%z.w.>.....
-...7..@GET /

14:17:51.407640 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.], seq 0:9, ack 1, win 4104, options
[nop,nop,TS val 766939831 ecr 932753216], length 9
E..=.%@.?.?(.....:W.P.f%z.w.>.....
-...7..@GET /

14:17:52.713721 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 1, win 4104, options [nop,nop,TS val
766941137 ecr 932753216], length 0
E..4.L@.?.....:W.P.f%.w.>.....\.....
-...7..@
14:17:53.914058 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.], seq 0:9, ack 1, win 4104, options
[nop,nop,TS val 766942335 ecr 932753216], length 9
E..=..N@.?.....:W.P.f%z.w.>.....'.....
-...7..@GET /

14:17:58.837755 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 2837, win 4051, options [nop,nop,TS
val 766947252 ecr 932765679], length 0
E..42.@.?..x.....:W.P.f%.w.R.....
-...7...
14:17:59.782728 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 5673, win 4007, options [nop,nop,TS
val 766948195 ecr 932765679], length 0
E..4..@.?..>.....:W.P.f%.w.f.....T.....
-..c7...
14:17:59.783120 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 5673, win 4096, options [nop,nop,TS
val 766948195 ecr 932765679], length 0
E..4. @.?.....:W.P.f%.w.f.....
-..c7...
14:18:00.701028 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 8509, win 4007, options [nop,nop,TS
val 766949110 ecr 932765679], length 0
E..4w @.?..O.....:W.P.f%.x z....w.....
-...7...
14:18:00.701428 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 8509, win 4096, options [nop,nop,TS
val 766949110 ecr 932765679], length 0
E..4V.@.?..pL.....:W.P.f%.x z....wT.....
-...7...
14:18:01.622806 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 11345, win 4007, options
[nop,nop,TS val 766950029 ecr 932766582], length 0
E..4.i@.?.....:W.P.f%.x.....e{.....
-...7..v
14:18:01.623010 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 11345, win 4096, options
[nop,nop,TS val 766950029 ecr 932766582], length 0
E..4..@.?..=.....:W.P.f%.x.....e".....
-...7..v
14:18:02.544587 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 14181, win 4007, options
[nop,nop,TS val 766950950 ecr 932767485], length 0
E..4..@.?.....:W.P.f%.x.....SG.....
-..&7...
14:18:02.544949 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [..], ack 14181, win 4096, options
[nop,nop,TS val 766950950 ecr 932767485], length 0
E..4..@.?..1r.....:W.P.f%.x.....R.....
-..&7...
```

## History

#1 - 06/18/2017 06:22 PM - sepherosa

Don't use fastforwarding, if you use firewall. BTW, I didn't find fastforwarding helps on any modern CPUs w/ medium amount of hosts on either side of the router.

On Mon, Jun 19, 2017 at 7:26 AM,

<[bugtracker-admin@leaf.dragonflybsd.org](mailto:bugtracker-admin@leaf.dragonflybsd.org)> wrote:

> Issue [#3045](#) has been reported by benjolitz.

>

> -----

> Bug [#3045](#): net.inet.ip.fastforwarding causes PF NAT to break/duplicate packets

> <http://bugs.dragonflybsd.org/issues/3045>

>

> \* Author: benjolitz

> \* Status: New

> \* Priority: Normal

> \* Assignee:

> \* Category: Networking

> \* Target version: Latest stable

> -----

> This issue took two days to track down.

>

> Here is the gist of the issue:

>

> If you set net.inet.ip.fastforwarding = 1

> AND set your pf.conf to:

>

> int\_if="int1"

> ext\_if="ext0"

> localnet=\$int\_if:network

> nat pass on \$ext\_if inet from \$localnet to any -> (\$ext\_if)

>

> Your tcpdump (for a telnet google.com 80 with `GET /` and a new line entered into the buffer) will show retransmits and unbelievably slow NAT behavior (we're talking about one blob of 120 bytes every 30 to 60 seconds)

>

> This bug is filed in hopes that google will pick up the following terms:

> pf nat slow, pf nat tcp duplicate

>

> If fastforwarding is disabled, PF will behave correctly.

>

> Here is a sample of it:

>

> 14:17:46.365163 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 3933726782, win 4104, options [nop,nop,TS val 766934808 ecr 932753216], length 0

> E..4?.@.?..2.....W.P.f%z.w.>.....

> -..7..@

> 14:17:47.171679 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 1, win 4104, options [nop,nop,TS val 766935613 ecr 932753216], length 0

> E..4vE@.?.P.....W.P.f%z.w.>.....

> -..=7..@

> 14:17:47.412131 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.] , seq 0:7, ack 1, win 4104, options [nop,nop,TS val 766935853 ecr 932753216], length 7

> E.;..@.?.....W.P.f%z.w.>....(.....

> -..-7..@GET /

>

> 14:17:47.556687 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.] , seq 0:7, ack 1, win 4104, options [nop,nop,TS val 766935998 ecr 932753216], length 7

> E.;;Z@.?.....W.P.f%z.w.>....'.....

> -..7..@GET /

>

> 14:17:47.708820 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.] , seq 0:7, ack 1, win 4104, options [nop,nop,TS val 766936150 ecr 932753216], length 7

> E.;;9.@.?.....W.P.f%z.w.>....'\_.....

> -..V7..@GET /

>

> 14:17:48.104957 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.] , seq 0:9, ack 1, win 4104, options [nop,nop,TS val 766936542 ecr 932753216], length 9

> E..=.@.?.A

> .....W.P.f%z.w.>.....

> -..7..@GET /

>

>

> 14:17:48.698278 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.] , seq 0:9, ack 1, win 4104, options [nop,nop,TS val 766937127 ecr 932753216], length 9

> E..=.H@.?.....W.P.f%z.w.>.....

> -..7..@GET /

>  
>  
> 14:17:48.700218 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 1, win 4104, options [nop,nop,TS val 766937129 ecr 932753216], length 0  
> E..4.@.?.....:W.P.f%..w.>.....  
> -..7..@  
> 14:17:49.668670 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.] , seq 0:9, ack 1, win 4104, options [nop,nop,TS val 766938095 ecr 932753216], length 9  
> E..=. @.?.c.....:W.P.f%z.w.>.....  
> -...7..@GET /  
>  
>  
> 14:17:51.407640 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.] , seq 0:9, ack 1, win 4104, options [nop,nop,TS val 766939831 ecr 932753216], length 9  
> E..=.%@.?.(.....:W.P.f%z.w.>.....  
> -...7..@GET /  
>  
>  
> 14:17:52.713721 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 1, win 4104, options [nop,nop,TS val 766941137 ecr 932753216], length 0  
> E..4.L@.?.....:W.P.f%..w.>.....\.....  
> -...7..@  
> 14:17:53.914058 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [P.] , seq 0:9, ack 1, win 4104, options [nop,nop,TS val 766942335 ecr 932753216], length 9  
> E..=.N@.?.....:W.P.f%z.w.>.....'  
> -...7..@GET /  
>  
>  
> 14:17:58.837755 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 2837, win 4051, options [nop,nop,TS val 766947252 ecr 932765679], length 0  
> E..42.@.?..x.....:W.P.f%..w.R.....  
> -...7...  
> 14:17:59.782728 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 5673, win 4007, options [nop,nop,TS val 766948195 ecr 932765679], length 0  
> E..4.@.?..>.....:W.P.f%..w.f.....T.....  
> -..c7...  
> 14:17:59.783120 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 5673, win 4096, options [nop,nop,TS val 766948195 ecr 932765679], length 0  
> E..4. @.?.....:W.P.f%..w.f.....  
> -..c7...  
> 14:18:00.701028 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 8509, win 4007, options [nop,nop,TS val 766949110 ecr 932765679], length 0  
> E..4w @.?..O.....:W.P.f%.x z....w.....  
> -...7...  
> 14:18:00.701428 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 8509, win 4096, options [nop,nop,TS val 766949110 ecr 932765679], length 0  
> E..4V.@.?..pL.....:W.P.f%.x z....wT.....  
> -...7...  
> 14:18:01.622806 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 11345, win 4007, options [nop,nop,TS val 766950029 ecr 932766582], length 0  
> E..4.i@.?.....:W.P.f%.x.....e{.....  
> -...7..v  
> 14:18:01.623010 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 11345, win 4096, options [nop,nop,TS val 766950029 ecr 932766582], length 0  
> E..4.@.?..=.....:W.P.f%.x.....e".....  
> -...7..v  
> 14:18:02.544587 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 14181, win 4007, options [nop,nop,TS val 766950950 ecr 932767485], length 0  
> E..4.@.?.....:W.P.f%.x.....SG.....  
> -..&7...  
> 14:18:02.544949 IP BenJolitz-Laptop.lan.50519 > lax17s05-in-f14.1e100.net.http: Flags [.] , ack 14181, win 4096, options [nop,nop,TS val 766950950 ecr 932767485], length 0  
> E..4.@.?..1r.....:W.P.f%.x.....R.....  
> -..&7...  
>  
>  
>  
> --  
> You have received this notification because you have either subscribed to it, or are involved in it.  
> To change your notification preferences, please click here: <http://bugs.dragonflybsd.org/my/account>

--  
Tomorrow Will Never Die

**#2 - 08/29/2017 10:14 AM - poige**

This is expected by nature of involved sub-system's parts.

If you're looking for kernel's alternative paths for routing there's PF's route-to as well.

**#3 - 06/19/2019 06:05 PM - liweitianux**

- *Status changed from New to Resolved*

This issue has been answered.

Sephe suggested not to use 'fastforwarding'.