# pfSense - Feature #2129

## TCP mss clamping for IPv6

01/21/2012 04:43 PM - Seth Mos

| | | | |
|---|---|---|---|
| **Status:** | Resolved | **Start date:** | 01/21/2012 |
| **Priority:** | Normal | **Due date:** | |
| **Assignee:** | | **% Done:** | 0% |
| **Category:** | Operating System | **Estimated time:** | 0.00 hour |
| **Target version:** | 2.2 | | |

**Description**

There is no tcp mss clamping for ipv6 and pf is not doing it either. Rumor has it that OpenBSD has it now.

Your system can not send or receive fragmented traffic over IPv6. The path between our system and your network has an MTU of 1480 bytes. The bottleneck is at IP address 2001:470:0:7d::2. The path between our system and your network does not appear to handle fragmented IPv6 traffic properly.

### History

**#1 - 01/22/2012 12:14 PM - JohnPoz _**

Where are you testing this exactly - my tests to http://test-ipv6.com/ show ok

Test IPv6 large packet
ok (0.204s) using ipv6
http://ipv6.test-ipv6.com/ip/?callback=?&size=1600&fill=xxx...xxx

Validates that IPv6 requests with large packets work. If this test times out, but other IPv6 tests work, it suggests that there may be PMTUD issues; possibly involving IP tunnels.

But when testing for edns getting changing results

```
dig +short rs.dns-oarc.net txt
rst.x3827.rs.dns-oarc.net.
rst.x3837.x3827.rs.dns-oarc.net.
rst.x3843.x3837.x3827.rs.dns-oarc.net.
"192.221.138.129 sent EDNS buffer size 4096"
"192.221.138.129 DNS reply size limit is at least 3843"
"Tested at 2012-01-22 17:14:36 UTC"
[2.1-DEVELOPMENT][root@pfsense.local.lan]/root(37): dig +short rs.dns-oarc.net txt
rst.x996.rs.dns-oarc.net.
rst.x1241.x996.rs.dns-oarc.net.
rst.x1294.x1241.x996.rs.dns-oarc.net.
"2001:470:1f10:b85::2 DNS reply size limit is at least 1294"
"2001:470:1f10:b85::2 sent EDNS buffer size 4096"
"Tested at 2012-01-22 17:14:39 UTC"
```

And showing that they are being blocked??

Jan 22 11:14:37 pfsense pf: 00:07:09.380498 rule 1/0(match): block in on re1: (tos 0x20, ttl 102, id 256, offset 0, flags [none], proto TCP (6), length 40)  125.210.214.76.6000 > 24.13.176.20.1433: Flags [S], cksum 0x282d (correct), seq 241958912, win 16384, length 0
Jan 22 11:14:37 pfsense pf: 00:00:47.294925 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::133 > 2001:470:1f10:b85::2: frag (0x0b81afc9:0|1432) 53 > 15027: 53398*- 1/82/82 rs.dns-oarc.net. CNAME rst.x4091.rs.dns-oarc.net. (1424)
Jan 22 11:14:37 pfsense pf: 00:00:00.000926 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::133 > 2001:470:1f10:b85::2: frag (0x0b81afc9:1432|1432)
Jan 22 11:14:37 pfsense pf: 00:00:00.000094 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1243) 2001:4f8:3:2bc:2::133 > 2001:470:1f10:b85::2: frag (0x0b81afc9:2864|1235)
Jan 22 11:14:37 pfsense pf: 00:00:00.010495 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::133 > 2001:470:1f10:b85::2: frag (0x5fd1cddd:0|1432) 53 > 15027: 53398*- 1/42/42 rs.dns-oarc.net. CNAME rst.x2031.rs.dns-oarc.net. (1424)
Jan 22 11:14:37 pfsense pf: 00:00:00.000929 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 615) 2001:4f8:3:2bc:2::133 > 2001:470:1f10:b85::2: frag (0x5fd1cddd:1432|607)
Jan 22 11:14:37 pfsense pf: 00:00:00.103623 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::133 > 2001:470:1f10:b85::2: frag (0x2e381314:0|1432) 53 > 50841: 37953- 0/82/82 (1424)
Jan 22 11:14:37 pfsense pf: 00:00:00.000931 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::133 > 2001:470:1f10:b85::2: frag (0x2e381314:1432|1432)

Jan 22 11:14:37 pfsense pf: 00:00:00.000037 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1223) 2001:4f8:3:2bc:2::133 > 2001:470:1f10:b85::2: frag (0x2e381314:2864|1215)
Jan 22 11:14:37 pfsense pf: 00:00:00.566702 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::134 > 2001:470:1f10:b85::2: frag (0x1ee1b7e9:0|1432) 53 > 32265: 26192*- 1/40/40 rst.x996.rs.dns-oarc.net. CNAME rst.x1946.x996.rs.dns-oarc.net. (1424)
Jan 22 11:14:37 pfsense pf: 00:00:00.000091 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 530) 2001:4f8:3:2bc:2::134 > 2001:470:1f10:b85::2: frag (0x1ee1b7e9:1432|522)
Jan 22 11:14:37 pfsense pf: 00:00:00.000068 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::134 > 2001:470:1f10:b85::2: frag (0x5eb5d4a9:0|1432) 53 > 32265: 26192*- 1/32/32 rst.x996.rs.dns-oarc.net. CNAME rst.x1570.x996.rs.dns-oarc.net. (1424)
Jan 22 11:14:37 pfsense pf: 00:00:00.000036 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 154) 2001:4f8:3:2bc:2::134 > 2001:470:1f10:b85::2: frag (0x5eb5d4a9:1432|146)
Jan 22 11:14:37 pfsense pf: 00:00:00.106994 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::134 > 2001:470:1f10:b85::2: frag (0x00ce2f00:0|1432) 53 > 54320: 35452- 0/82/82 (1424)
Jan 22 11:14:37 pfsense pf: 00:00:00.000921 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::134 > 2001:470:1f10:b85::2: frag (0x00ce2f00:1432|1432)
Jan 22 11:14:37 pfsense pf: 00:00:00.000097 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1229) 2001:4f8:3:2bc:2::134 > 2001:470:1f10:b85::2: frag (0x00ce2f00:2864|1221)
Jan 22 11:14:37 pfsense pf: 00:00:00.820849 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::134 > 2001:470:1f10:b85::2: frag (0x2059d1e2:0|1432) 53 > 19719: 3896- 0/82/82 (1424)
Jan 22 11:14:37 pfsense pf: 00:00:00.000961 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::134 > 2001:470:1f10:b85::2: frag (0x2059d1e2:1432|1432)
Jan 22 11:14:37 pfsense pf: 00:00:00.000105 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1229) 2001:4f8:3:2bc:2::134 > 2001:470:1f10:b85::2: frag (0x2059d1e2:2864|1221)
Jan 22 11:14:37 pfsense pf: 00:00:00.632964 rule 3/0(match): block in on gif0: (hlim 57, next-header Fragment (44) payload length: 1440) 2001:4f8:3:2bc:2::134 > 2001:470:1f10:b85::2: frag (0x6b0e12eb:0|1432) 53 > 40404: 11530- 0/82/82 (1424)

Seeing this traffic on my tunnel endpoint, vs my lans routed /64 ip which is in 1f11 block vs that 1f10 block.  I been playing with what rules I cold put into allow the blocked traffic.  But no matter what rules I put in it seems to just not like the fragment stuff?

**#2 - 01/22/2012 12:30 PM - JohnPoz _**

Ok I ran a test at ICSI Netalyzr, and ran into the same thing

IPv6 Path MTU (?): Warning
Your system can not send or receive fragmented traffic over IPv6. The path between our system and your network has an MTU of 1480 bytes. The bottleneck is at IP address 2001:470:0:6e::2. The path between our system and your network does not appear to handle fragmented IPv6 traffic properly.

If you look up that ip I show this
; <<>> DiG 9.6.2-P2 <<>> -x 2001:470:0:6e::2
PTR gige-gbge0.tserv9.chi1.ipv6.he.net.

Seems to point to a problem in He network??

#### #3 - 01/22/2012 03:58 PM - Seth Mos

the problem is that pf currently does not handle ipv6 fragments. the other case is IPsec tunnels and other VPN solutions where that is really just neccesary.

Also, the IPv6 network is going to fall to the same pitfalls that made path mtu on ipv4 a issue. It's wishful thinking we won't need it. Cisco doesn't have it, but I already ran into a case where it's needed.

#### #4 - 03/21/2012 02:42 PM - Chris Buechler

*- Target version changed from 8 to 2.1*

#### #5 - 10/26/2012 02:56 PM - Chris Buechler

*- Tracker changed from Bug to Feature*

*- Subject changed from No TCP mss clamping for IPv6 to TCP mss clamping for IPv6*

*- Target version changed from 2.1 to 2.2*

#### #6 - 05/13/2014 09:40 PM - Chris Buechler

*- Target version deleted (2.2)*

questionable whether this is necessary. Definitely not a priority for 2.2

#### #7 - 11/08/2014 07:38 AM - Doktor Notor

Chris Buechler wrote:

> questionable whether this is necessary. Definitely not a priority for 2.2

If you question whether it's neccessary, perhaps you should read these:

http://robert.penz.name/971/google-services-seems-to-be-down-if-youre-accessing-them-via-an-ipv6-tunnel-providers/
http://lists.cluenet.de/pipermail/ipv6-ops/2014-November/010255.html

Also, the GUI does it wrong. E.g., when you set up MSS for a GIF IPv6 tunnel, it takes the value and sets minus 40 (which is wrong for IPv6, kindly see RFC2460 -> 8.3) at https://tools.ietf.org/html/rfc2460#page-28. Why not just take what the user inputs and fix the description goes beyond me.

#### #8 - 11/08/2014 03:04 PM - Chris Buechler

*- Status changed from New to Resolved*

*- Target version set to 2.2*

*- Affected Documentation 0 added*

MTU in RA and properly-functioning PMTUD do indeed make it questionable as to whether it's necessary. But MSS clamping does work on v6 in 2.2.

The MSS clamping field is the way it is because people understand the total end result packet size more so than trying to figure out payload+headers. That does need some adjustment for v6, but it still works as described (clamping to 40 bytes less than value entered). We're well aware of how big an

IPv6 header is, Notor.

#### #9 - 11/09/2014 02:20 PM - Doktor Notor

Ok, so people understand better that the input value is not taken as input value but subtracted by some (incorrect) number. Sorry, this makes about zero sense. Either you can remove the MSS settings altogether and set it to MTU - 40 for IPv4 and MTU - 60 for IPv6 automagically, or let people set what they **really** want to set instead.