

IPsec tcp session stalling (me too) ...

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2005-10/0292.html>

From: Matthew Grooms (*mgrooms_at_shrew.net*)

Date: 10/22/05

Date: Sat, 22 Oct 2005 13:33:27 -0500

To: volker@vwsoft.com

Volker,

I have noticed the same problem. In my case, it only seems to happen when the traffic is being forwarded across interfaces and pf or ipfw is enabled. I use purely IPSEC so I would agree that GRE isn't the problem. This behavior is 100% reproducible for me. If traffic is forwarded from the host providing the ESP protection or if the firewall package is disabled, the problem goes away.

Just some data points. I don't recall seeing this ever happen on 4.x + ipfw. I experienced this on early 5.x + ipfw, late 5.x + pf and 6.x + pf. I believe the ipfw versions I tested were prior to the pfil hooks conversion.

For example ...

NODE 1 sftp client

NODE 2 sftp server

IPSEC policy requires ESP protection from NODE 1 or VPN A to NODE 2

NODE 1 ----- VPN A ===== VPN B ----- NODE 2

- 1) NODE 1 <-> NODE 2 sftp via IPSEC pf enabled, traffic stalls
- 2) NODE 1 <-> NODE 2 sftp via IPSEC pf disabled, no problems
- 3) VPN A <-> NODE 2 sftp via IPSEC pf enabled, no problems

NOTE : TCP protocol is irrelevant. Haven't tried UDP.

tcpdump from VPN A internal interface in the (1) case ...

```
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on xl1, link-type EN10MB (Ethernet), capture size 96 bytes
12:47:02.673536 IP 10.22.200.21.1223 > 10.20.10.141.ssh: S
929298208:929298208(0) win 16384 <mss 1460,nop,nop,sackOK>
12:47:02.714815 IP 10.20.10.141.ssh > 10.22.200.21.1223: S
4112220692:4112220692(0) ack 929298209 win 5840 <mss 1460,nop,nop,sackOK>
12:47:02.715005 IP 10.22.200.21.1223 > 10.20.10.141.ssh: . ack 1 win 17520
```

freebsd-net: IPsec tcp session stalling (me too) ...

```
12:47:04.390884 IP 10.20.10.141.ssh > 10.22.200.21.nerv: P
4052971969:4052972505(536) ack 992811526 win 6432
12:47:04.391082 IP 10.22.200.21.nerv > 10.20.10.141.ssh: . ack 536 win 16319
12:47:06.312301 IP 10.20.10.141.ssh > 10.22.200.21.1223: S
4112220692:4112220692(0) ack 929298209 win 5840 <mss 1460,nop,nop,sackOK>
12:47:06.312418 IP 10.22.200.21.1223 > 10.20.10.141.ssh: . ack 1 win 17520
12:47:12.311382 IP 10.20.10.141.ssh > 10.22.200.21.1223: S
4112220692:4112220692(0) ack 929298209 win 5840 <mss 1460,nop,nop,sackOK>
12:47:12.311492 IP 10.22.200.21.1223 > 10.20.10.141.ssh: . ack 1 win 17520
12:47:12.349209 IP 10.20.10.141.ssh > 10.22.200.21.1223: P 1:26(25) ack
1 win 5840
12:47:12.349463 IP 10.22.200.21.1223 > 10.20.10.141.ssh: P 1:38(37) ack
26 win 17495
12:47:15.342481 IP 10.20.10.141.ssh > 10.22.200.21.1223: P 1:26(25) ack
1 win 5840
12:47:15.342592 IP 10.22.200.21.1223 > 10.20.10.141.ssh: . ack 26 win 17495
12:47:15.530728 IP 10.22.200.21.1223 > 10.20.10.141.ssh: P 1:446(445)
ack 26 win 17495
12:47:22.093189 IP 10.22.200.21.1223 > 10.20.10.141.ssh: P 1:446(445)
ack 26 win 17495
12:47:22.140791 IP 10.20.10.141.ssh > 10.22.200.21.1223: . ack 446 win 6432
12:47:22.141931 IP 10.20.10.141.ssh > 10.22.200.21.1223: P 26:666(640)
ack 446 win 6432
12:47:22.142232 IP 10.22.200.21.1223 > 10.20.10.141.ssh: P 446:462(16)
ack 666 win 16855
12:47:28.136807 IP 10.20.10.141.ssh > 10.22.200.21.1223: P 26:666(640)
ack 446 win 6432
12:47:28.137018 IP 10.22.200.21.1223 > 10.20.10.141.ssh: . ack 666 win 16855
12:47:35.218016 IP 10.22.200.21.1223 > 10.20.10.141.ssh: P 446:462(16)
ack 666 win 16855
```

tcpdump from VPN A internal interface in the (2) case ...

```
root@hole# tcpdump -i xl1 src or dst 10.20.10.141
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on xl1, link-type EN10MB (Ethernet), capture size 96 bytes
12:58:36.411839 IP 10.22.200.21.1228 > 10.20.10.141.ssh: S
3459356125:3459356125(0) win 16384 <mss 1460,nop,nop,sackOK>
12:58:36.452844 IP 10.20.10.141.ssh > 10.22.200.21.1228: S
548348284:548348284(0) ack 3459356126 win 5840 <mss 1460,nop,nop,sackOK>
12:58:36.452954 IP 10.22.200.21.1228 > 10.20.10.141.ssh: . ack 1 win 17520
12:58:40.048592 IP 10.20.10.141.ssh > 10.22.200.21.1228: S
548348284:548348284(0) ack 3459356126 win 5840 <mss 1460,nop,nop,sackOK>
12:58:40.048693 IP 10.22.200.21.1228 > 10.20.10.141.ssh: . ack 1 win 17520
12:58:40.086207 IP 10.20.10.141.ssh > 10.22.200.21.1228: P 1:26(25) ack
1 win 5840
12:58:40.086452 IP 10.22.200.21.1228 > 10.20.10.141.ssh: P 1:38(37) ack
26 win 17495
12:58:40.120738 IP 10.20.10.141.ssh > 10.22.200.21.1228: . ack 38 win 5840
12:58:40.120915 IP 10.22.200.21.1228 > 10.20.10.141.ssh: P 38:446(408)
ack 26 win 17495
```

IPsec tcp session stalling (me too) ...

freebsd-net: IPSec tcp session stalling (me too) ...

12:58:40.124873 IP 10.20.10.141.ssh > 10.22.200.21.1228: P 26:666(640)
ack 38 win 5840
12:58:40.199077 IP 10.20.10.141.ssh > 10.22.200.21.1228: . ack 446 win 6432
12:58:40.199198 IP 10.22.200.21.1228 > 10.20.10.141.ssh: P 446:462(16)
ack 666 win 16855
12:58:40.237038 IP 10.20.10.141.ssh > 10.22.200.21.1228: . ack 462 win 6432
12:58:40.245737 IP 10.20.10.141.ssh > 10.22.200.21.1228: P 666:1202(536)
ack 462 win 6432
12:58:40.317230 IP 10.22.200.21.1228 > 10.20.10.141.ssh: P 462:990(528)
ack 1202 win 16319
12:58:40.390865 IP 10.20.10.141.ssh > 10.22.200.21.1228: . ack 990 win 7504
12:58:40.544484 IP 10.20.10.141.ssh > 10.22.200.21.1228: P
1202:2242(1040) ack 990 win 7504
12:58:40.594600 IP 10.22.200.21.1228 > 10.20.10.141.ssh: P 990:1006(16)
ack 2242 win 17520
12:58:40.628544 IP 10.20.10.141.ssh > 10.22.200.21.1228: . ack 1006 win 7504
12:58:40.628660 IP 10.22.200.21.1228 > 10.20.10.141.ssh: P 1006:1054(48)
ack 2242 win 17520
12:58:40.663963 IP 10.20.10.141.ssh > 10.22.200.21.1228: . ack 1054 win 7504
12:58:40.664821 IP 10.20.10.141.ssh > 10.22.200.21.1228: P 2242:2290(48)
ack 1054 win 7504
12:58:40.665010 IP 10.22.200.21.1228 > 10.20.10.141.ssh: P 1054:1118(64)
ack 2290 win 17472
12:58:40.723227 IP 10.20.10.141.ssh > 10.22.200.21.1228: P 2290:2370(80)
ack 1118 win 7504
12:58:40.723421 IP 10.22.200.21.1228 > 10.20.10.141.ssh: P 1118:1198(80)
ack 2370 win 17392
12:58:40.783036 IP 10.20.10.141.ssh > 10.22.200.21.1228: P 2370:2402(32)
ack 1198 win 7504
12:58:40.783226 IP 10.22.200.21.1228 > 10.20.10.141.ssh: P 1198:1262(64)
ack 2402 win 17360
12:58:40.817686 IP 10.20.10.141.ssh > 10.22.200.21.1228: P 2402:2450(48)
ack 1262 win 7504
12:58:40.817926 IP 10.22.200.21.1228 > 10.20.10.141.ssh: P 1262:1342(80)
ack 2450 win 17312
12:58:40.856225 IP 10.20.10.141.ssh > 10.22.200.21.1228: P 2450:2482(32)
ack 1342 win 7504
12:58:40.856392 IP 10.22.200.21.1228 > 10.20.10.141.ssh: P 1342:1390(48)
ack 2482 win 17280
12:58:40.903573 IP 10.20.10.141.ssh > 10.22.200.21.1228: P 2482:2562(80)
ack 1390 win 7504
12:58:40.903921 IP 10.20.10.141.ssh > 10.22.200.21.1228: P
2562:2690(128) ack 1390 win 7504
12:58:40.904061 IP 10.22.200.21.1228 > 10.20.10.141.ssh: . ack 2690 win
17072
12:58:40.942770 IP 10.20.10.141.ssh > 10.22.200.21.1228: P 2690:2738(48)
ack 1390 win 7504
12:58:41.088509 IP 10.22.200.21.1228 > 10.20.10.141.ssh: . ack 2738 win
17024

freebsd-net: IPSec tcp session stalling (me too) ...

Sorry in advance for not posting as a reply to the original message. I don't subscribe to the list. Just wanted to substantiate Volkers findings.

Hope this helps,

Matthew Grooms

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"