pfSense - Bug #7801

UDP fragments received over IPsec tunnel are not properly reassembled and forwarded

08/22/2017 09:17 PM - Chris Linstruth

Status:	New	**Start date:**	08/22/2017
Priority:	Normal	**Due date:**	
Assignee:	Luiz Souza	**% Done:**	0%
Category:	IPsec	**Estimated time:**	0.00 hour
Target version:	2.5.next		
Affected Version:	2.3.4_1	**Affected Architecture:**	

Description

```
I used this to generate the large UDP datagrams on Host J1:
nping —udp —source-port 5060 —dest-port 5060 —data-length 2000 172.31.200.100
or
nping —udp —source-port 5060 —dest-port 5060 —data-length 2000 172.31.202.100

With pf enabled on both H and J

Into J LAN
01:12:22.571417 IP 172.31.201.100.5060 > 172.31.200.100.5060: UDP, length 2000
01:12:22.571425 IP 172.31.201.100 > 172.31.200.100: ip-proto-17
01:12:23.599657 IP 172.31.201.100.5060 > 172.31.200.100.5060: UDP, length 2000
01:12:23.599666 IP 172.31.201.100 > 172.31.200.100: ip-proto-17
01:12:24.603630 IP 172.31.201.100.5060 > 172.31.200.100.5060: UDP, length 2000
01:12:24.603639 IP 172.31.201.100 > 172.31.200.100: ip-proto-17

Into IPsec on J
01:13:29.451155 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000
01:13:30.463517 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000
01:13:31.482213 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000

Out of IPsec on H
01:15:00.971973 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000
01:15:01.997529 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000
01:15:03.029548 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000

Out H LAN
01:16:09.025710 IP 172.31.201.100.5060 > 172.31.202.100.5060: UDP, length 2000
01:16:09.025715 IP 172.31.201.100 > 172.31.202.100: ip-proto-17
01:16:10.047232 IP 172.31.201.100.5060 > 172.31.202.100.5060: UDP, length 2000
01:16:10.047239 IP 172.31.201.100 > 172.31.202.100: ip-proto-17
01:16:11.066202 IP 172.31.201.100.5060 > 172.31.202.100.5060: UDP, length 2000
01:16:11.066209 IP 172.31.201.100 > 172.31.202.100: ip-proto-17

(Note that ICMP unreachables were returned by 172.31.202.100 since there is nothing
listening on udp/5060. They were received by nping. This verifies good end-to-end
connectivity.)

Disable pf on pfSense J to avoid fragment reassembly / scrubbing

With pf enabled on H and disabled on J

Into J LAN
```

```
01:21:32.108916 IP 172.31.201.100.5060 > 172.31.200.100.5060: UDP, length 2000
01:21:32.108982 IP 172.31.201.100 > 172.31.200.100: ip-proto-17
01:21:33.132073 IP 172.31.201.100.5060 > 172.31.200.100.5060: UDP, length 2000
01:21:33.132135 IP 172.31.201.100 > 172.31.200.100: ip-proto-17
01:21:34.153074 IP 172.31.201.100.5060 > 172.31.200.100.5060: UDP, length 2000
01:21:34.153133 IP 172.31.201.100 > 172.31.200.100: ip-proto-17

Into IPsec on J
01:22:13.844357 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000
01:22:13.844409 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100 > 172.31.200.100: ip-p
roto-17
01:22:14.853298 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000
01:22:14.853390 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100 > 172.31.200.100: ip-p
roto-17
01:22:15.857681 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000
01:22:15.857729 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100 > 172.31.200.100: ip-p
roto-17

Out of IPsec on H
01:23:08.108706 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000
01:23:08.108729 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100 > 172.31.200.100: ip-p
roto-17
01:23:09.109471 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000
01:23:09.109493 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100 > 172.31.200.100: ip-p
roto-17
01:23:10.115393 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100.5060 > 172.31.200.100.
5060: UDP, length 2000
01:23:10.115415 (authentic,confidential): SPI 0xc99a1e36: IP 172.31.201.100 > 172.31.200.100: ip-p
roto-17

Out H LAN
[No packets captured]

Note that removing the BINAT on IPsec on H changes this behavior. These captures
are the same (Replacing destination with 172.31.202.100 because no more BINAT):

Into J LAN
Into IPsec on J
Out of IPsec on H

But the following is captured on H LAN
01:41:02.519981 IP 172.31.201.100.5060 > 172.31.202.100.5060: UDP, length 2000
01:41:03.524742 IP 172.31.201.100.5060 > 172.31.202.100.5060: UDP, length 2000
01:41:04.528723 IP 172.31.201.100.5060 > 172.31.202.100.5060: UDP, length 2000

Note: Not fragmented according to the interface’s mtu 1500. This is similar to
what has been reported on #7779 for OpenVPN.

Upgraded J to 2.4-RC - Same behavior
Upgraded H to 2.4-RC - Same behavior
```

History

#1 - 09/11/2017 12:08 AM - Jim Thompson

- Assignee set to Jim Pingle

#2 - 09/11/2017 03:56 PM - Renato Botelho

- Target version changed from 2.4.0 to 2.4.1

#3 - 10/12/2017 09:46 AM - Jim Pingle

- Target version changed from 2.4.1 to 2.4.2

Moving target to 2.4.2 as we need 2.4.1 sooner than anticipated.

#4 - 10/23/2017 12:39 PM - Jim Pingle

- Target version changed from 2.4.2 to 2.4.3

#5 - 02/01/2018 12:46 PM - Steve Beaver

- Assignee changed from Jim Pingle to Luiz Souza

- Target version changed from 2.4.3 to 2.4.4

++target_version

#6 - 08/09/2018 09:00 AM - Franciszek Koltuniuk

Hi,
I have a similar issue with fragmented packets send/received over IPsec tunnel.
Finally, I manage to update */tmp/rules.debug* and make traffic to be passed as expected:

- add *scrub* rule that apply scrubbing on a traffic from networks behind IPsec vpn:

  scrub from <from_ipsec_network> to any no-df fragment reassemble

- add *scrub* rult that disable scrubbing on a traffic from a local interface to the network behind IPsec vpn:

  no scrub on $INTERFACE to <from_ipsec_network>

So the only one thing is missing: pfsense web interface that allows to setup ipsec in the same way..

#7 - 08/14/2018 01:52 PM - Steve Beaver

- Target version changed from 2.4.4 to 48

#8 - 12/14/2018 02:00 AM - Andi Admin

Any chance to get fixed soon? This bug actually prevent our VPN from being usable for VoIP which uses UDP and in some cases need fragments.

#9 - 01/04/2019 07:46 AM - Next Next

Hi, I also have a similar issue with fragmented packets and IPsec tunnels (noticed with ICMP traffic).
Incoming fragmented packets (either from 'LAN to IPsec' or 'IPsec to LAN') are forwarded unfragmented (packets larger then interface MTU).
Disabling scrubbing globally makes the problem go away, but this doesn't seem to be a good solution as it affects much more than only bug-related traffic.

Is there any info on what's causing this? Is it PFsense/config related or maybe something in BSD itself ?

#10 - 01/11/2019 07:06 PM - Gabriel Latour

Hi, I have been waiting a year for that fix, for us, it's RDS sessions that disconnects randomly when using UDP over IPSEC.

#11 - 03/12/2019 10:54 AM - Jim Pingle

- Target version changed from 48 to 2.5.0

#12 - 08/14/2019 09:09 AM - Jim Pingle

See also: #9184, #7837

#13 - 11/05/2020 07:30 AM - Steve Beaver

- Target version changed from 2.5.0 to 2.5.next

Files

pfSense+VPN copy 2.png	45.5 KB	08/23/2017	Chris Linstruth