strongSwan - Issue #3580

encapsulation and packets not routing into tunnel problems

26.09.2020 22:10 - Tom jung

Status:	Feedback		
Priority:	Normal		
Assignee:			
Category:			
Affected version:	5.7.2	**Resolution:**	

Description

IPsec SA: only UDP encapsulation is supported
unable to install inbound and outbound IPsec SA (SAD) in kernel
Cannot get any packets to route into the tunnel

encapsulation and packets not routing into tunnel problems

New implementation with 2 problems:
1st
IPsec SA: only UDP encapsulation is supported.
unable to install inbound and outbound IPsec SA (SAD) in kernel.

Sep 26 15:00:14 ip-10-0-58-73.ec2.internal strongswan[9625]: 16[CFG] selected proposal: ESP:AES_CBC_256
Sep 26 15:00:14 ip-10-0-58-73.ec2.internal strongswan[9625]: 16[ESP] IPsec SA: only UDP encapsulation
Sep 26 15:00:14 ip-10-0-58-73.ec2.internal strongswan[9625]: 16[ESP] failed to create SAD entry
Sep 26 15:00:14 ip-10-0-58-73.ec2.internal strongswan[9625]: 16[ESP] IPsec SA: only UDP encapsulation
Sep 26 15:00:14 ip-10-0-58-73.ec2.internal strongswan[9625]: 16[ESP] failed to create SAD entry

2nd
Also, cannot get any packets to route into the tunnel from our test server.

From test server, a ping or telnet to a port on the other end of the VPN shows traffic not going into tunnel but instead going to strongSwan default gateway 10.0.56.1.

tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on eth0, link-type EN10MB (Ethernet), capture size 262144 bytes
15:06:12.572350 IP 10.0.56.157 > 158.73.182.48: ICMP echo request, id 1, seq 1391, length 40
15:06:12.572383 IP 10.0.58.73 > 10.0.56.157: ICMP redirect 158.73.182.48 to host 10.0.56.1, length 68

History

#1 - 28.09.2020 15:00 - Tobias Brunner

- Status changed from New to Feedback

- Priority changed from Urgent to Normal

Why are you using the kernel-libipsec plugin? And if you do, why is UDP encapsulation not forced (the plugin usually does that)?

#2 - 28.09.2020 17:11 - Tom jung

I didn't make that choice for kernel-libipsec plugin, what am I suppose to be using?

#3 - 28.09.2020 17:46 - Tobias Brunner

> I didn't make that choice for kernel-libipsec plugin, what am I suppose to be using?

How come you are using it then? On Linux, you probably want to use the *kernel-netlink* plugin.

#4 - 28.09.2020 18:01 - Tom jung

So are you saying SS the install makes kernel-libipsec plugin the default, and instead I need to change that in the ipsec.conf like the Host-to-Host Tunnels section indicates?

#5 - 28.09.2020 18:04 - Tobias Brunner

> So are you saying SS the install makes kernel-libipsec plugin the default

It's definitely not the default (see e.g. Autoconf or PluginList).

> and instead I need to change that in the ipsec.conf like the Host-to-Host Tunnels section indicates?

Sorry, no idea what you are referring to. But if you wonder how to enable/disable plugins, see PluginLoad.

#6 - 29.09.2020 20:26 - Tom jung

Uninstalled kernel-libipsec plugin and encapsulation error has disappeared, Phase 2 has come up. We still cannot route any traffic into tunnel, it keeps going to SS default gateway. We are trying to do Source NATing and have tried many combinations with no luck.
Any advice or do you need more details?

#7 - 30.09.2020 12:45 - Tobias Brunner

> Any advice or do you need more details?

Definitely need more details to provide any specific advice. Generally, just make sure the negotiated IPsec policies match the actual traffic (i.e. after the SNAT).

#8 - 30.09.2020 15:50 - Tom jung

- File VPN diagram.JPG added

- File etc-sysconfig-iptables.txt added

- File etc-iptables.conf.txt added

- File etc-rc.local.txt added

- File etc-strongswan-ipsec.conf.txt added

- File etc-sysconfig-network-scripts-route-eth0.txt added

- File etc-sysctl.conf.txt added

#9 - 30.09.2020 15:55 - Tom jung

I've attached the 6 configs (excluding secrets) we've modified in their current state of many, and a diagram. We are currently using a test server 10.0.56.157 instead of the two production servers in the diagram to route traffic into the tunnel. Hope this helps clarify what we're trying to do.

#10 - 30.09.2020 15:57 - Tobias Brunner

ipsec statusall?

#11 - 30.09.2020 15:58 - Tom jung

- File ip route show table all.txt added

7th file I forgot to include, route table.

#12 - 30.09.2020 16:03 - Tom jung

- File systemctl status strongswan.txt added

This version doesn't support ipsec statusall. Heres systemctl status strongswan

#13 - 30.09.2020 16:16 - Tobias Brunner

> This version doesn't support ipsec statusall.

Then it's probably strongswan statusall.

> Heres systemctl status strongswan

Why would you think that helps?

#14 - 30.09.2020 16:20 - Tom jung

- File strongswan statusall.txt added

Didn't know of that strongswan statusall command. Here you go!

#15 - 30.09.2020 16:21 - Tobias Brunner

Please run that again after the connection is up.

#16 - 30.09.2020 16:25 - Tom jung

- File strongswan statusall.txt added

Updated strongswan statusall after systemctl restart strongswan

#17 - 30.09.2020 16:28 - Tobias Brunner

> Updated strongswan statusall after systemctl restart strongswan

Looks fine (except for the horrible algorithms and protocol version). So if your packets come from 7.7.20.5 (after the SNAT) and are addressed to 158.73.182.48, they should get tunneled.

#18 - 30.09.2020 16:31 - Tom jung

Here is the strongswan.conf that it says "no files found matching '/etc/strongswan/strongswan.conf'". This is after editing trying to NOT use kernel-libipsec so I need to go back an clean it up since we just removed the kernel-libipsec plugin.

#19 - 30.09.2020 16:32 - Tom jung

- File etc-strongswan-strongswan.conf.txt added

Attachment strongswan.conf

#20 - 30.09.2020 16:45 - Tom jung

- File strongswan statusall.txt added

Updated strongswan statusall after putting back original /etc/strongswan/strongswan.conf

#21 - 30.09.2020 16:56 - Tom jung

- File traceroute from test server to host 158.73.182.48 at other end of VPN .txt added

Traceroute just keeps showing SS sending traffic to its default gateway and not into tunnel (attached).

#22 - 30.09.2020 17:01 - Tobias Brunner

That's because your NAT rule is incorrect. You used -d instead of -s.

#23 - 30.09.2020 17:31 - Tom jung

Is there a command to see if there are packets going into the tunnel? The remote side doesn't have a test server like me so they won't respond to my connection attempts like normal.

#24 - 30.09.2020 17:36 - Tom jung

This what my iptables looks like now.

```
[tjung@ip-10-0-58-73 ~]$ sudo iptables -t nat -L
Chain PREROUTING (policy ACCEPT)
target     prot opt source               destination

Chain INPUT (policy ACCEPT)
target     prot opt source               destination

Chain OUTPUT (policy ACCEPT)
```

target prot opt source destination

Chain POSTROUTING (policy ACCEPT)
target prot opt source destination
SNAT all -- ip-10-0-56-157.ec2.internal anywhere to:7.7.20.5

#25 - 30.09.2020 17:44 - Tobias Brunner

> Is there a command to see if there are packets going into the tunnel?

statusall shows the packet/byte counters and for iptables you can use -v to see counters for each rule. You should also see outbound ESP packets in tcpdump.

#26 - 30.09.2020 18:36 - Tom jung

I never see ESP packets in my tcpdump (example attached earlier). While I see my pkts count go up when pinging from test server, remote end ASA says they never see packets in the tunnel.

Chain POSTROUTING (policy ACCEPT 867 packets, 64852 bytes)
pkts bytes target prot opt in out source destination
5 396 SNAT all -- any any ip-10-0-56-157.ec2.internal anywhere to:7.7.20.5

#27 - 30.09.2020 18:42 - Tobias Brunner

> I never see ESP packets in my tcpdump (example attached earlier)

Obviously, as the NAT rule was wrong. If it is correct now, i.e. the packets' addresses match the IPsec policy, you should see some.

> While I see my pkts count go up when pinging from test server, remote end ASA says they never see packets in the tunnel.
>
> Chain POSTROUTING (policy ACCEPT 867 packets, 64852 bytes)
> pkts bytes target prot opt in out source destination
> 5 396 SNAT all -- any any ip-10-0-56-157.ec2.internal anywhere to:7.7.20.5

That's only the NAT rule. What about the packet counters for the IPsec SAs?

#28 - 30.09.2020 18:52 - Tom jung

- File strongswan statusall.txt added

Updated statusall attached. After a SS restart, I see my count go up to FDS-to-McK_CB{1}: 3DES_CBC/HMAC_SHA1_96, 0 bytes_i, 240 bytes_o (4 pkts, 7s ago). Remote end ASA still says they never see packets.

#29 - 30.09.2020 18:57 - Tobias Brunner

> Remote end ASA still says they never see packets.

Maybe ESP is blocked by a firewall between the two hosts. You could try forcing UDP encapsulation via *forceencaps=yes*.

#30 - 30.09.2020 19:00 - Tom jung

We're not using the firewall, but in the etc-strongswan-ipsec.conf.txt from earlier you'll see we already have forceencaps=yes

#31 - 30.09.2020 19:04 - Tobias Brunner

> We're not using the firewall

There is no firewall whatsoever between your IKE/IPsec endpoints?

> but in the etc-strongswan-ipsec.conf.txt from earlier you'll see we already have forceencaps=yes

That's not reflected by the output of statusall, which says ESP and not ESP in UDP. Maybe the peer does not support NAT-Traversal (or has it disabled).

#32 - 30.09.2020 19:06 - Tom jung

Our peer?

#33 - 30.09.2020 19:07 - Tobias Brunner

> Our peer?

The other IKE implementation.

#34 - 30.09.2020 19:22 - Tom jung

- File strongswan statusall.txt added

Updated statusall attached. Now my tcpdump no longer shows this like earlier 09:51:48.646169 IP 10.0.58.73 > 10.0.56.157: ICMP 10.0.58.73 udp port netbios-ns unreachable, length 86

#35 - 30.09.2020 19:48 - Tom jung

Here is my feedback from the remote ASA guy. NAT-Traversal is not enabled. As long as your VPN device (SS) is using a public IP address or is receiving a static NAT to a public IP address, NAT-T does not come into play.

#36 - 30.09.2020 20:16 - Tom jung

- File Fortigate NAT traversal.JPG added

On our current production VPN using our Fortigate for this same business partner, NAT Traversal is Enable on our side (attached), so I would think SS should be doing the same. If so how do I do that?

#37 - 30.09.2020 20:38 - Tom jung

According to this, it looks like SS is doing NAT-T just like our Fortigate is. Do you agree that NAT-T is not the problem?
IKEv1
Before strongSwan 5.0.0, NAT discovery and traversal for IKEv1 had to be enabled by setting nat_traversal=yes in the config setup section of ipsec.conf. Otherwise, strongSwan 4.x's IKEv1 pluto daemon would not accept incoming IKE packets with a UDP source port different from 500. ***Since 5.0.0 IKEv1 traffic is handled by the charon daemon, which supports NAT traversal according to RFC 3947 (and some of its early drafts) without having to enable it explicitly (it can't be disabled either, though).***

#38 - 30.09.2020 20:50 - Tom jung

- File swanctl -l.txt added

Attached swanctl -l. Are you able to join a screen share troubleshooting with me and ASA guy this afternoon?

#39 - 01.10.2020 09:24 - Tobias Brunner

> As long as your VPN device (SS) is using a public IP address or is receiving a static NAT to a public IP address, NAT-T does not come into play.

The whole point of **forcing** encapsulation is to use NAT-T to fake a NAT situation even if there is no NAT in order to enable UDP encapsulation.

> Are you able to join a screen share troubleshooting with me and ASA guy this afternoon?

No.

Since you are sending ESP packets now, it's either a problem on the way to the other end (e.g. a firewall blocking ESP) or at the other end (could be a firewall or routing problem).

#40 - 01.10.2020 15:33 - Tom jung

Help me understand, since uninstalling the kernel-libipsec plugin and the encapsulation error going away, the ipsec0 interface is gone (below). How does the traffic get into the tunnel then?

[tjung@ip-10-0-58-73 ~]$ sudo netstat -i
Kernel Interface table

```
Iface      MTU    RX-OK RX-ERR RX-DRP RX-OVR    TX-OK TX-ERR TX-DRP TX-OVR Flg
eth0      9001   473270      0      0 0       507323      0      0      0 BMRU
lo       65536        0      0      0 0            0      0      0      0 LRU
```

Where are these OUT packets going?
sudo swanctl -l
out 1c938b96, 60 bytes, 1 packets, 3s ago

#41 - 01.10.2020 16:07 - Tobias Brunner

> How does the traffic get into the tunnel then?

IPsec policies (see ip xfrm policy) that match the traffic direct them to the IPsec SAs (ip xfrm state).

> Where are these OUT packets going?

They get processed by the outbound IPsec SA (the traffic counters there are actually those of that SA) and then sent as ESP packets.

#42 - 01.10.2020 17:02 - Tom jung

How can I verify NAT traversal is actually working?

#43 - 01.10.2020 17:18 - Tobias Brunner

> How can I verify NAT traversal is actually working?

Read the log, or check the status output (as mentioned above).

#44 - 01.10.2020 18:20 - Tom jung

- File ip xfrm state.txt added

- File ip xfrm policy.txt added

Attached are the 2 you mentioned. I don't think the xfrm state is correct, it has the remote ASA public but SS private IP (src 10.0.58.73 dst 139.177.139.10).
Shouldn't src be the SS public IP which is 52.86.64.169?

#45 - 01.10.2020 22:51 - Tom jung

- File fw (6).pcap added

I had my remote ASA guy run a Pcap (attached), it's all ISAKMP, he never sees any ESP packets.

#46 - 02.10.2020 10:03 - Tobias Brunner

> I don't think the xfrm state is correct, it has the remote ASA public but SS private IP (src 10.0.58.73 dst 139.177.139.10).
> Shouldn't src be the SS public IP which is 52.86.64.169?

Depends on the IP addresses and routes on your system and the strongSwan config.

> I had my remote ASA guy run a Pcap (attached), it's all ISAKMP, he never sees any ESP packets.

So try to find out where the ESP packets are dropped.

Files

etc-sysconfig-iptables.txt	494 Bytes	30.09.2020	Tom jung
VPN diagram.JPG	43.7 KB	30.09.2020	Tom jung
etc-iptables.conf.txt	91 Bytes	30.09.2020	Tom jung
etc-rc.local.txt	619 Bytes	30.09.2020	Tom jung
etc-strongswan-ipsec.conf.txt	1006 Bytes	30.09.2020	Tom jung

etc-sysconfig-network-scripts-route-eth0.txt	166 Bytes	30.09.2020	Tom jung
etc-sysctl.conf.txt	236 Bytes	30.09.2020	Tom jung
ip route show table all.txt	1.59 KB	30.09.2020	Tom jung
systemctl status strongswan.txt	1.55 KB	30.09.2020	Tom jung
strongswan statusall.txt	1.19 KB	30.09.2020	Tom jung
strongswan statusall.txt	1.85 KB	30.09.2020	Tom jung
etc-strongswan-strongswan.conf.txt	632 Bytes	30.09.2020	Tom jung
strongswan statusall.txt	1.75 KB	30.09.2020	Tom jung
traceroute from test server to host 158.73.182.48 at other end of VPN .txt	1.13 KB	30.09.2020	Tom jung
strongswan statusall.txt	3.51 KB	30.09.2020	Tom jung
strongswan statusall.txt	2.49 KB	30.09.2020	Tom jung
Fortigate NAT traversal.JPG	70.9 KB	30.09.2020	Tom jung
swanctl -l.txt	8.81 KB	30.09.2020	Tom jung
ip xfrm state.txt	770 Bytes	01.10.2020	Tom jung
ip xfrm policy.txt	1.09 KB	01.10.2020	Tom jung
fw (6).pcap	15.1 KB	01.10.2020	Tom jung