I was always curious how the IPSec session looks like after PAT translation. As we discovered in IPSec basics: IPSec through NAT article, IPSec must use some NAT-avoiding mechanism to work through NAT/PAT. I have to say (for those who aren't IPSec fan) that most IPSec connections are made through the NAT (at least most Remote Access VPN connections). So that is a common case when IPSec session encapsulated in udp packets (in case of NAT-T).

Let's see what happens with one (first) IPSec session before and after PAT.

Legend: 192.168.100.* are internal hosts (IPSec-initiators); 192.0.2.2 is a PAT address (internet address)

First IPSec session encapsulated in udp packets (NAT-T) with PAT

Sess. No.	Source IP/port before PAT	Source IP/port after PAT
1		

192.168.100.11:500 (IKE negot. messages)
192.168.100.11:4500 (IPSec "stream")

192.0.2.2:500 (IKE negot. messages)
192.0.2.2:4500 (IPSec "stream")

Seems to me PAT translation was pretty trivial one. All source ports are unchanged and as far as I remember it's default behavior of cisco's NAT algorithm.

But often there are multiple IPsec-initiator behind the same PAT. For example, our Network Security Team used to open three or even five sessions to the same customer's VPN-device (implementation phase, your know :)). Let's see how does it goes when there are multiple sessions.

Multiple IPSec sessions encapsulated in udp packets (NAT-T) with PAT

Sess. No.	Source IP/port before PAT	Source IP/port after PAT
1	192.168.100.11:500 (IKE negot. messages)	192.0.2.2:500 (IKE negot. messages)
	192.168.100.11:4500 (IPSec "stream")	192.0.2.2:4500 (IPSec "stream")
2	192.168.100.12:500 (IKE negot. messages)	192.0.2.2:1 (IKE negot. messages)
	192.168.100.12:4500 (IPSec "stream")	192.0.2.2:1024 (IPSec "stream")
3	192.168.100.13:500 (IKE negot. messages)	192.0.2.2:2 (IKE negot. messages)
	192.168.100.13:4500 (IPSec "stream")	192.0.2.2:1025 (IPSec "stream")

Looks curious, isn't it? :)

Since the udp:500 and udp:4500 ports are taken by the first session all subsequent sessions must use something different. As we can suppose cisco's NAT algorithm translate all "duplicated" source ports below or equal 1023 to 1,2,3, et cetera and all "duplicated" source port between 1024 and 65535 to 1024, 1025, 1026 et cetera.

You can say "And what is the usage case for all this stuff?". You know, some people know that IPSec works through udp:500/4500 ports and use the ACL to control IPSec sessions as an extra security measure. Once I had a call from one of mine customers about "weird IPSec issue" and I was lucky enough to knew all the answers ;).