strongSwan - Issue #951

p2 SA are not deleted on FreeBSD

11.05.2015 19:46 - Florian Apolloner

Status:	Closed		
Priority:	Normal		
Assignee:	Tobias Brunner		
Category:	kernel		
Affected version:	5.3.0	**Resolution:**	No change required

Description

I am having a weird problem where old SA are not deleted, but their expiry counter increases:

```
Security Associations (3 up, 0 connecting):
     con3000[12]: ESTABLISHED 28 minutes ago, YYY.YYY.YYY.YYY[YYY.YYY.YYY.YYY]...ZZZ.ZZZ.ZZZ.ZZZ[1
72.31.255.1]
     con3000[12]: IKEv1 SPIs: 2c2a79ee0522ea76_i* 7d61779e65e5152a_r, pre-shared key reauthenticat
ion in 18 minutes
     con3000[12]: IKE proposal: 3DES_CBC/HMAC_MD5_96/PRF_HMAC_MD5/MODP_1024
     con3000{19}:  REKEYED, TUNNEL, reqid 3, expires in 1 second
     con3000{19}:   172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0
     con3000{22}:  INSTALLED, TUNNEL, reqid 3, ESP SPIs: cc682556_i 7ae0af2a_o
     con3000{22}:  3DES_CBC/HMAC_MD5_96, 3287 bytes_i (17 pkts, 691s ago), 6656 bytes_o (20 pkts,
643s ago), rekeying in 41 seconds
     con3000{22}:   172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0

Security Associations (3 up, 0 connecting):
     con3000[12]: ESTABLISHED 29 minutes ago, YYY.YYY.YYY.YYY[YYY.YYY.YYY.YYY]...ZZZ.ZZZ.ZZZ.ZZZ[1
72.31.255.1]
     con3000[12]: IKEv1 SPIs: 2c2a79ee0522ea76_i* 7d61779e65e5152a_r, pre-shared key reauthenticat
ion in 18 minutes
     con3000[12]: IKE proposal: 3DES_CBC/HMAC_MD5_96/PRF_HMAC_MD5/MODP_1024
     con3000{19}:  REKEYED, TUNNEL, reqid 3, expires in 6 seconds
     con3000{19}:   172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0
     con3000{22}:  INSTALLED, TUNNEL, reqid 3, ESP SPIs: cc682556_i 7ae0af2a_o
     con3000{22}:  3DES_CBC/HMAC_MD5_96, 3287 bytes_i (17 pkts, 698s ago), 6656 bytes_o (20 pkts,
650s ago), rekeying in 34 seconds
     con3000{22}:   172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0

Security Associations (3 up, 0 connecting):
     con3000[12]: ESTABLISHED 29 minutes ago, YYY.YYY.YYY.YYY[YYY.YYY.YYY.YYY]...ZZZ.ZZZ.ZZZ.ZZZ[1
72.31.255.1]
     con3000[12]: IKEv1 SPIs: 2c2a79ee0522ea76_i* 7d61779e65e5152a_r, pre-shared key reauthenticat
ion in 17 minutes
     con3000[12]: IKE proposal: 3DES_CBC/HMAC_MD5_96/PRF_HMAC_MD5/MODP_1024
     con3000{19}:  REKEYED, TUNNEL, reqid 3, expires in 40 seconds
     con3000{19}:   172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0
     con3000{22}:  INSTALLED, TUNNEL, reqid 3, ESP SPIs: cc682556_i 7ae0af2a_o
     con3000{22}:  3DES_CBC/HMAC_MD5_96, 6490 bytes_i (33 pkts, 732s ago), 11352 bytes_o (36 pkts,
 5s ago), rekeying in 0 seconds
     con3000{22}:   172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0

Security Associations (3 up, 0 connecting):
     con3000[12]: ESTABLISHED 29 minutes ago, YYY.YYY.YYY.YYY[YYY.YYY.YYY.YYY]...ZZZ.ZZZ.ZZZ.ZZZ[1
72.31.255.1]
     con3000[12]: IKEv1 SPIs: 2c2a79ee0522ea76_i* 7d61779e65e5152a_r, pre-shared key reauthenticat
ion in 17 minutes
     con3000[12]: IKE proposal: 3DES_CBC/HMAC_MD5_96/PRF_HMAC_MD5/MODP_1024
     con3000{19}:  REKEYED, TUNNEL, reqid 3, expires in 41 seconds
     con3000{19}:   172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0
     con3000{22}:  REKEYED, TUNNEL, reqid 3, expires in 16 minutes
     con3000{22}:   172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0
     con3000{23}:  INSTALLED, TUNNEL, reqid 3, ESP SPIs: ca8a6594_i 57f4a68b_o
```

```
     con3000{23}:  3DES_CBC/HMAC_MD5_96, 0 bytes_i, 0 bytes_o, rekeying in 14 minutes
     con3000{23}:   172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0
```

As you can see the old SA approaches zero and then goes up again! Same goes for the "new old" SA after the rekey.

My config:

```
conn con3000
        reqid = 3
        fragmentation = yes
        keyexchange = ikev1
        reauth = yes
        forceencaps = no
        mobike = no
        rekey = yes
        installpolicy = yes
        type = tunnel
        dpdaction = none
        auto = route
        left = YYY.YYY.YYY.YYY
        right = ZZZ.ZZZ.ZZZ.ZZZ
        leftid = YYY.YYY.YYY.YYY
        ikelifetime = 3600s
        lifetime = 1740s
        ike = 3des-md5-modp1024!
        esp = 3des-md5!
        leftauth = psk
        rightauth = psk
        rightid = 172.31.255.1
        aggressive = no
        rightsubnet = XXX.XXX.XXX.XXX/24
        leftsubnet = 172.17.40.0/24|172.22.1.0/24
```

The logs I did enable show the following:

```
May 11 19:28:31 gw01 charon: 13[ENC] <con3000|12> generating QUICK_MODE request 1118644536 [ HASH
SA No ID ID ]
May 11 19:28:31 gw01 charon: 13[NET] <con3000|12> sending packet: from YYY.YYY.YYY.YYY[500] to ZZZ
.ZZZ.ZZZ.ZZZ[500] (172 bytes)
May 11 19:28:31 gw01 charon: 13[NET] <con3000|12> received packet: from ZZZ.ZZZ.ZZZ.ZZZ[500] to YY
Y.YYY.YYY.YYY[500] (156 bytes)
May 11 19:28:31 gw01 charon: 13[ENC] <con3000|12> parsed QUICK_MODE response 1118644536 [ HASH SA
No ID ID ]
May 11 19:28:31 gw01 charon: 13[IKE] <con3000|12> CHILD_SA con3000{23} established with SPIs ca8a6
594_i 57f4a68b_o and TS 172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0
May 11 19:28:31 gw01 charon: 13[IKE] <con3000|12> CHILD_SA con3000{23} established with SPIs ca8a6
594_i 57f4a68b_o and TS 172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0
May 11 19:28:31 gw01 charon: 13[ENC] <con3000|12> generating QUICK_MODE request 1118644536 [ HASH
]
May 11 19:28:31 gw01 charon: 13[NET] <con3000|12> sending packet: from YYY.YYY.YYY.YYY[500] to ZZZ
.ZZZ.ZZZ.ZZZ[500] (52 bytes)
```

Related issues:

Related to Issue #1103: Stuck with rekeying active	**Closed**	**10.09.2015**
Has duplicate Issue #1250: REKEYING problem between strongSwan and MikroTik r...	**Closed**	**31.12.2015**

History

#1 - 12.05.2015 14:48 - Tobias Brunner

- Tracker changed from Bug to Issue

- Status changed from New to Feedback

```
    leftsubnet = 172.17.40.0/24|172.22.1.0/24
```

I've never seen this syntax, nor this:

```
con3000{23}:   172.17.40.0/24|172.22.1.0/24 === XXX.XXX.XXX.XXX/24|/0
```

What version are you using? On what platform? Which what patches?

#2 - 12.05.2015 14:50 - Florian Apolloner

Strongswan 5.3.0 on PfSense 2.2.2, no idea how or what they patched :/

Downstream bugreport: https://redmine.pfsense.org/issues/4686

#3 - 12.05.2015 21:56 - Florian Apolloner

Regarding the odd syntax for leftsubnet: This tells pfsense (according to the webinterface), that while my network is 172.22.1.0/24 it should be netmaped (in the iptables sense, or binat for freebsd users) to 172.17.40.0/24 for the other side.

EDIT:// After looking at the relevant strongswan patches in pfsense (They are behind a CLA to sign, so I'll not publish them here for now, just to be safe), it seems as if this patch is really just there to provide a "nicer" display in status/statusall -- for every other purpose the pipe is replaced with \0. That being said, I fully understand that you might not want to invest any time in this issue given downstream patching. Either way, I'd appreciate if you could still take a look and see if it would be possible somehow that the expiry time rises after going through zero (even though I cannot guarantee that I did not miss anything in the pfsense patches :/)

#4 - 13.05.2015 19:17 - Tobias Brunner

- File freebsd-send-hard-expire.patch added

Thanks for the additional information. It looks like this is actually a general FreeBSD issue, i.e. is not caused by the pfSense patches.

strongSwan depends on the kernel to send messages when soft or hard lifetimes of IPsec SAs expire to trigger the appropriate action (rekey or delete). Additionally, for IKEv1 old IPsec SAs are not immediately deleted after they have been rekeyed. Instead, they are just kept installed until they expire naturally (i.e. via hard lifetime).

Unfortunately, the FreeBSD kernel apparently does not send an SADB_EXPIRE message to the keying daemon when the hard lifetime of an SA is reached. The kernel just silently removes it, so the daemon never learns that the SA expired. The lifetime used in ipsec statusall is the precomputed absolute expiry time and simply logged as absolute difference to the current time. So the reported time will increase after the SA expired.

I also saw that SAs are only removed by the kernel if a soft lifetime has been set, as the SA otherwise never reaches SADB_SASTATE_DYING. So *rekeyfuzz=0%* does cause the outbound SA to never expire as we only set a soft lifetime on the inbound SA in that case to avoid triggering multiple soft expires at the same time (this wouldn't be that much of a problem if we'd actually receive the hard expire for the inbound SA as we'd then delete both SAs anyway).

According to RFC 2367, section 3.1.8 FreeBSD should probably send an SADB_EXPIRE when the hard lifetime is reached:

> The operating system kernel is responsible for tracking SA
> expirations for security protocols that are implemented inside the
> kernel. If the soft limit **or hard limit** of a Security Association
> has expired for a security protocol implemented inside the kernel,
> then the kernel MUST issue an SADB_EXPIRE message to all key socket
> listeners.

It later continues (this is also how we differentiate between soft and hard expire when we receive an SADB_EXPIRE):

> If a HARD lifetime extension is included, it indicates that the HARD
> lifetime expired. This means the association MAY be deleted already
> from the SADB. If a SOFT lifetime extension is included, it indicates
> that the SOFT lifetime expired.

But the RFC also specifies:

> The messaging behavior of the SADB_EXPIRE message is:
>
> The kernel sends an SADB_EXPIRE message to all listeners when
> the soft limit of a security association has been expired.

Here for some reason only the soft limit is mentioned, but that might be an oversight.

It's not much of an issue to fix this in the kernel, the attached patch actually does exactly that. I'll probably submit that later to the FreeBSD bug tracker.

A workaround would be to patch the kernel-pfkey plugin so it would schedule a hard expire in userland on FreeBSD (but I'd rather fix this in the kernel to be honest).

#5 - 13.05.2015 21:05 - Ermal Luçi

Just for the record, the interesting part is that the subnet part after |(pipe) is used on nat situations to create two P2 in the kernel with lan side the original subnet and peer side with translated one.
This way you can apply pure nat firewall rules on enc(4) interface on FreeBSD rather than other ways of doing it.

#6 - 16.05.2015 20:25 - Florian Apolloner

Tobias Brunner wrote:

> Thanks for the additional information. It looks like this is actually a general FreeBSD issue, i.e. is not caused by the pfSense patches.

Is this an actual issue or just cosmetic? Ie after a certain time nothing goes through the tunnel anymore and I am wondering if that is caused by that behaviour…

> It's not much of an issue to fix this in the kernel, the attached patch actually does exactly that. I'll probably submit that later to the FreeBSD bug tracker.

Got a link to the ticket for me?

#7 - 18.05.2015 16:38 - Tobias Brunner

> Thanks for the additional information. It looks like this is actually a general FreeBSD issue, i.e. is not caused by the pfSense patches.
>
> Is this an actual issue or just cosmetic? Ie after a certain time nothing goes through the tunnel anymore and I am wondering if that is caused by that behaviour…

It should mainly be cosmetic if *rekeyfuzz* is not set to *0%* (the default is *100%*). However, some additional memory is required in the daemon for each expired SA.

Besides the *rekeyuzz* problem I explained in my other comment another issue is that soft expires are only sent to the daemon if the SAs have actually been used since they were installed, this could eventually result in no SA being established once they have expired (using *auto=route* might counter this problem). I addressed this with another patch (see bug reports below).

> It's not much of an issue to fix this in the kernel, the attached patch actually does exactly that. I'll probably submit that later to the FreeBSD bug tracker.
>
> Got a link to the ticket for me?

I filed two bug reports today: 200282 and 200283

#8 - 08.07.2015 11:16 - Tobias Brunner

- Subject changed from p2 SA are not deleted to p2 SA are not deleted on FreeBSD

- Category set to kernel

- Status changed from Feedback to Closed

- Assignee set to Tobias Brunner

- Resolution set to No change required

The FreeBSD kernel fixes have been merged a while ago.

#9 - 10.09.2015 15:47 - Tobias Brunner

- Related to Issue #1103: Stuck with rekeying active added

#10 - 18.01.2016 11:34 - Tobias Brunner

- Has duplicate Issue #1250: REKEYING problem between strongSwan and MikroTik router (v6.3.3) added

Files