strongSwan - Issue #1288

Special SAD inserted by charon when traffic is not up

01.02.2016 03:36 - heidi rao

Status:	Closed		
Priority:	Normal		
Assignee:	Tobias Brunner		
Category:	kernel		
Affected version:	5.3.5	**Resolution:**	No change required

Description

Hi,
The config on host is as below:

```
# ipsec.conf
# FlexiPlatform: IPsec configuration file

config setup
    uniqueids=no
    charondebug="knl 2,enc 0,net 0,ike 2,cfg 2,mgr 2,chd 2"
    conn %default
    auto=route
    keyexchange=ikev2
    reauth=no
    ca RULEABC-01~VPNABC-1
    cacert="/etc/ipsec.d/cacerts/"
    conn RULEABC-01~VPNABC-1
    rekeymargin=6
    rekeyfuzz=100%
    keyexchange=ikev1
    left=34.2.2.200
    right=34.2.2.10
    leftsubnet=34.2.2.0/24
    rightsubnet=34.2.2.0/24
    authby=rsasig
    leftcert="/etc/ipsec.d/certs/fpccert.pem"
    leftid=34.2.2.200
    rightid=%any
    ike=aes128-sha1-modp768!
    esp=aes128-sha1!
    type=tunnel
    ikelifetime=500s
    keylife=500s
    mobike=no
    auto=route
    reauth=no
```

The peer is not up, execute command "ping 34.2.2.101 -I 34.2.2.201", a special SAD appeared as below:

```
[root@24F-VFPC-002 ~]# ip xfrm state
src 34.2.2.200 dst 34.2.2.10
        proto esp spi 0x00000000 reqid 1 mode tunnel
        replay-window 0
        sel src 34.2.2.201/32 dst 34.2.2.101/32 proto icmp type 8 code 0
```

If I restart charon by command "ipsec restart", the SAD is still there, however, after a while it disappeared.
Does this SAD inserted by charon? If yes, which file and function in source code it's mapped to? Thanks!

Heidi

History

#1 - 01.02.2016 09:35 - Tobias Brunner

- Subject changed from Specail SAD inserted by charon when traffic is not up to Special SAD inserted by charon when traffic is not up

- Description updated

- Category changed from charon to kernel

- Status changed from New to Feedback

> Does this SAD inserted by charon?

No, the kernel creates these states to track matches for trap policies (*auto=route*). When traffic matches such a policy the kernel creates a temporary state and sends acquires to listening daemons, which will attempt to negotiate actual SAs. These states time out after a configurable amount of time (*charon.plugins.kernel-netlink.xfrm_acq_expires* or sysctl net.core.xfrm_acq_expires). The plugin sets this value to 165 seconds by default, which equals the default retransmission timeout.

Because current versions of strongSwan don't flush the kernel states when terminating (earlier versions did something similar to ip xfrm state flush) and these states are not tracked by the daemon they survive the termination.

#2 - 02.02.2016 12:00 - heidi rao

Hi,
Thanks!
Is there any API to clear this SA? I check the netlink interface, it use "struct xfrm_usersa_id" to notify kernel, it seems this data structure can't delete this SA.

Heidi

#3 - 02.02.2016 15:50 - Tobias Brunner

> Is there any API to clear this SA? I check the netlink interface, it use "struct xfrm_usersa_id" to notify kernel, it seems this data structure can't delete this SA.

No it can't. The only option is to either install and delete a matching SA or to flush all SAs via XFRM_MSG_FLUSHSA.

#4 - 23.05.2016 15:41 - Tobias Brunner

- Status changed from Feedback to Closed

- Assignee set to Tobias Brunner

- Resolution set to No change required