strongSwan - Feature #1243

Add support for overlapping trap policies

22.12.2015 10:24 - heidi rao

Status:	Closed	**Start date:**	22.12.2015
Priority:	Normal	**Due date:**	
Assignee:	Tobias Brunner	**Estimated time:**	0.00 hour
Category:	kernel-interface		
Target version:	5.5.2		
Resolution:	Fixed		

Description

Hi,
I configured two connections as below:

```
conn RULEABC-0~VPNABC-0
        rekeymargin=600
        rekeyfuzz=100%
        keyexchange=ikev2
        right=35.6.6.201
        left=35.6.6.101
        rightsubnet=35.6.6.202/32
        leftsubnet=35.6.6.102/32
        rightprotoport=1
        leftprotoport=1
        authby=secret
        leftid=35.6.6.101
        rightid=%any
        ike=3des-sha1-modp2048!
        esp=3des-sha1!
        type=tunnel
        ikelifetime=6000s
        keylife=6000s
        mobike=no
        auto=route
        reauth=no

conn RULEABC-1~VPNABC-1
        rekeymargin=600
        rekeyfuzz=100%
        keyexchange=ikev2
        right=35.6.6.205
        left=35.6.6.105
        rightsubnet=35.6.6.0/24
        leftsubnet=35.6.6.0/24
        leftprotoport=1
        rightprotoport=1
        authby=secret
        leftid=35.6.6.105
        rightid=%any
        ike=3des-sha1-modp2048!
        esp=aes128-sha1!
        type=tunnel
        ikelifetime=6000s
        keylife=6000s
        mobike=no
        auto=route
        reauth=no
```

After "ipsec restart", the SPD is shown as below, we can see the SPD with 32 mask has higher priority than the one with 24 mask.

```
src 35.6.6.0/24 dst 35.6.6.0/24 proto icmp
        dir fwd priority 5954 ptype main
        tmpl src 35.6.6.205 dst 35.6.6.105
                proto esp reqid 2 mode tunnel
src 35.6.6.0/24 dst 35.6.6.0/24 proto icmp
        dir in priority 5954 ptype main
        tmpl src 35.6.6.205 dst 35.6.6.105
                proto esp reqid 2 mode tunnel
src 35.6.6.0/24 dst 35.6.6.0/24 proto icmp
        dir out *priority 5954* ptype main
        tmpl src 35.6.6.105 dst 35.6.6.205
                proto esp reqid 2 mode tunnel
src 35.6.6.202/32 dst 35.6.6.102/32 proto icmp
        dir fwd priority 5890 ptype main
        tmpl src 35.6.6.201 dst 35.6.6.101
                proto esp reqid 1 mode tunnel
src 35.6.6.202/32 dst 35.6.6.102/32 proto icmp
        dir in priority 5890 ptype main
        tmpl src 35.6.6.201 dst 35.6.6.101
                proto esp reqid 1 mode tunnel
src 35.6.6.102/32 dst 35.6.6.202/32 proto icmp
        dir out *priority 5890* ptype main
        tmpl src 35.6.6.101 dst 35.6.6.201
                proto esp reqid 1 mode tunnel
```

After I execute command like "ping 35.6.6.206 -I 35.6.6.106", the SPD with 24 mask is updated, and has higher priority.

```
SPD
src 35.6.6.0/24 dst 35.6.6.0/24 proto icmp
        dir fwd priority 2882 ptype main
        tmpl src 35.6.6.205 dst 35.6.6.105
                proto esp reqid 2 mode tunnel
src 35.6.6.0/24 dst 35.6.6.0/24 proto icmp
        dir in priority 2882 ptype main
        tmpl src 35.6.6.205 dst 35.6.6.105
                proto esp reqid 2 mode tunnel
src 35.6.6.0/24 dst 35.6.6.0/24 proto icmp
        dir out *priority 2882* ptype main
        tmpl src 35.6.6.105 dst 35.6.6.205
                proto esp reqid 2 mode tunnel
src 35.6.6.202/32 dst 35.6.6.102/32 proto icmp
        dir fwd priority 5890 ptype main
        tmpl src 35.6.6.201 dst 35.6.6.101
                proto esp reqid 1 mode tunnel
src 35.6.6.202/32 dst 35.6.6.102/32 proto icmp
        dir in priority 5890 ptype main
        tmpl src 35.6.6.201 dst 35.6.6.101
                proto esp reqid 1 mode tunnel
src 35.6.6.102/32 dst 35.6.6.202/32 proto icmp
        dir out *priority 5890* ptype main
        tmpl src 35.6.6.101 dst 35.6.6.201

SAD
src 35.6.6.105 dst 35.6.6.205
        proto esp spi 0xcf26ac9d reqid 2 mode tunnel
        replay-window 32 flag 20
        auth hmac(sha1) 0x45ec41dad2ea4bca545fe099c908d8cc96a35b65
        enc cbc(aes) 0x815d078bdb65eb45e55126a3ed5d1944
src 35.6.6.205 dst 35.6.6.105
        proto esp spi 0xc369b792 reqid 2 mode tunnel
        replay-window 32 flag 20
        auth hmac(sha1) 0x4d148e787194074feacbd9439f138213a6dcc8fd
        enc cbc(aes) 0x1137ce60be2416d6c76c1b4b38ba1041
```

Then I execute "ping 35.6.6.202 -I 35.6.6.102", it always match the SPD with 24 mask, the traffic is sent to 35.6.6.205, not 35.6.6.201.

Why charon insert the trap policy with lower priority after the connection is routed, and updated after the child_sa setup? Thanks!

Heidi

Related issues:

Has duplicate Issue #2125: Priority of SPD will be updated after SAD created	**Closed**	**28.09.2016**

Associated revisions

Revision 0e9d6c46 - 08.02.2017 10:36 - Tobias Brunner

kernel-netlink: Use the same priority range for trap and regular policies

While trap and regular policies now often look the same (mainly because reqids are kept constant) trap policies still need to have a lower priority than regular policies to handle unroute/route correctly if e.g. IPComp is used or the mode changes. But if we use a completely different priority range that's lower than that of regular policies it is not possible to install overlapping trap policies. By differentiating trap from regular policies via the priority's LSB this issue is avoided while still maintaining the proper ordering of trap and regular policies.

Fixes #1243.

History

#1 - 23.12.2015 14:40 - Tobias Brunner

- Status changed from New to Feedback

The priorities are assigned because the Linux kernel manages policies in a linear list ordered by the assigned priority. So we assign priorities based on the traffic selector (subnet size etc.) and the type of the policy. The charon IKE daemon always installed trap policies with lower priority than the actual IPsec policies. I don't know the actual reason, but it might have been due to some early limitations (e.g. the priority calculation did originally not include the remote subnet) or perhaps because the old pluto daemon did it too (not sure if it did).

The priorities are currently also used to keep track of the policies. Since the Linux kernel can't handle duplicate policies we have to make sure to install the right one. Which is less of a problem since 5.3.0 because reqids don't change anymore for matching policies. So because trap and regular policies now pretty often look the same (an exception might be the combination with IPComp, or if transport mode switches to tunnel mode), we could perhaps modify this a bit and use the same priorities but still internally order routed policies after regular policies. But it is also only a problem with overlapping policies, as seen here, or due to narrowing as seen in #1245, which is not that common (and could possibly be avoided depending on the config).

#2 - 29.09.2016 10:06 - Tobias Brunner

- Has duplicate Issue #2125: Priority of SPD will be updated after SAD created added

#3 - 10.10.2016 12:56 - Xiaoqiang Fu

- File 0001-kernel-netlink-separate-routed-and-default-policy-wi.patch added

Hi,
I am focus on Issue #2125.
There is one question in v5.5.0:
1. Why is POLICY_PRIORITY_ROUTED and POLICY_PRIORITY_DEFAULT separated?

```
    switch (prio)
    {
        case POLICY_PRIORITY_FALLBACK:
            priority += PRIO_BASE;
            /* fall-through to next case */
        case POLICY_PRIORITY_ROUTED:
            priority += PRIO_BASE;
            /* fall-through to next case */
        case POLICY_PRIORITY_DEFAULT:
            priority += PRIO_BASE;
            /* fall-through to next case */
        case POLICY_PRIORITY_PASS:
            break;
    }
```

If POLICY_PRIORITY_ROUTED and POLICY_PRIORITY_DEFAULT have the same priority base, then this issue is resolved.

2. Another solution is using priority least bit to separate POLICY_PRIORITY_ROUTED and POLICY_PRIORITY_DEFAULT.
Could you please give you comments about the attached solution?

I will submit it if suitable.

#4 - 11.10.2016 15:28 - Tobias Brunner

- Tracker changed from Issue to Feature

- Subject changed from Priority of SPD will be updated after SAD created to Add support for overlapping trap policies

> 1. Why is POLICY_PRIORITY_ROUTED and POLICY_PRIORITY_DEFAULT separated?

Did you read what I wrote above? Internally they have to be ordered differently.

> If POLICY_PRIORITY_ROUTED and POLICY_PRIORITY_DEFAULT have the same priority base, then this issue is resolved.

No, it's not that easy. If you e.g. have a routed connection that uses IPComp and after establishing it unroute and then route it again the installed policies won't be correct if the priorities are the same.

> 2. Another solution is using priority least bit to separate POLICY_PRIORITY_ROUTED and POLICY_PRIORITY_DEFAULT.

Yes, that might work. I pushed a couple of commits to the *1243-trap-prio* branch.

#5 - 08.02.2017 10:37 - Tobias Brunner

- Category set to kernel-interface

- Status changed from Feedback to Closed

- Assignee set to Tobias Brunner

- Target version set to 5.5.2

- Resolution set to Fixed

Files

0001-kernel-netlink-separate-routed-and-default-policy-wi.patch	1.27 KB	10.10.2016	Xiaoqiang Fu