strongSwan - Issue #1517

IPv6 IPsec connection (transport mode) is not re-established if initiator side is restarted

15.06.2016 15:02 - Jiri Zendulka

Status:	Closed		
Priority:	Normal		
Assignee:	Tobias Brunner		
Category:	configuration		
Affected version:	5.4.0	**Resolution:**	No change required

Description

Initiator side ----ipv6/transport mode --- Responder side

The connection is succesfully established after first start. If I restart initiator side the connection is not automatically re-established. IPv4 works well - connection is re-established.

The responder side:

```
conn ipsec1
 left=fd07:7::ee
 right=%any
 leftauth=psk
 rightauth=psk
 leftfirewall=yes
 keyexchange=ikev2
 ikelifetime=3600
 keylife=3600
 rekeymargin=540
 rekeyfuzz=100%
 keyingtries=%forever
 type=transport
 ike=aes128-sha256-modp3072,aes128-sha1-modp2048,3des-sha1-modp1536
 esp=aes128-sha256,aes128-sha1,3des-sha1
 auto=add
```

```
Status of IKE charon daemon (weakSwan 5.4.0, Linux 3.12.10+, armv7l):
  uptime: 11 minutes, since Jun 15 12:32:46 2016
  malloc: sbrk 532480, mmap 0, used 130192, free 402288
  worker threads: 11 of 16 idle, 5/0/0/0 working, job queue: 0/0/0/0, scheduled: 3
  loaded plugins: charon nonce pem openssl kernel-netlink socket-default stroke updown
Listening IP addresses:
  192.168.1.1
  192.168.7.101
  fd07:7::ee
Connections:
      ipsec1:  fd07:7::ee...%any  IKEv2
      ipsec1:   local:  [fd07:7::ee] uses pre-shared key authentication
      ipsec1:   remote: uses pre-shared key authentication
      ipsec1:   child:  dynamic === dynamic TRANSPORT
Security Associations (1 up, 1 connecting):
   (unnamed)[13]: CONNECTING, fd07:7::ee[%any]...fd07:7::e5[%any]
   (unnamed)[13]: IKEv2 SPIs: fbff01842639b504_i 8cf55f3995b5cd0e_r*
   (unnamed)[13]: IKE proposal: AES_CBC_128/HMAC_SHA2_256_128/PRF_HMAC_SHA2_256/MODP_3072
   (unnamed)[13]: Tasks passive: IKE_CERT_PRE IKE_AUTH IKE_CERT_POST IKE_CONFIG CHILD_CREATE IKE_A
UTH_LIFETIME IKE_MOBIKE
      ipsec1[1]: ESTABLISHED 11 minutes ago, fd07:7::ee[fd07:7::ee]...fd07:7::e5[fd07:7::e5]
      ipsec1[1]: IKEv2 SPIs: ed19c995a7466941_i 94b3ad98fc2ee2d0_r*, pre-shared key reauthenticati
on in 36 minutes
      ipsec1[1]: IKE proposal: AES_CBC_128/HMAC_SHA2_256_128/PRF_HMAC_SHA2_256/MODP_3072
      ipsec1{1}:  INSTALLED, TRANSPORT, reqid 1, ESP SPIs: cf9f5228_i c1737c79_o
      ipsec1{1}:  AES_CBC_128/HMAC_SHA2_256_128, 3392 bytes_i (53 pkts, 637s ago), 44160 bytes_o (
690 pkts, 0s ago), rekeying in 36 minutes
      ipsec1{1}:   fd07:7::ee/128 === fd07:7::e5/128
```

```
2016-06-15 12:48:26 charon: 12[IKE] fd07:7::e5 is initiating an IKE_SA
2016-06-15 12:48:26 charon: 12[ENC] generating IKE_SA_INIT response 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) N(HASH_ALG) N(MULT_AUTH) ]
2016-06-15 12:48:26 charon: 12[NET] sending packet: from fd07:7::ee[500] to fd07:7::e5[500] (584 bytes)
2016-06-15 12:48:30 charon: 15[NET] received packet: from fd07:7::e5[500] to fd07:7::ee[500] (1280 bytes)
2016-06-15 12:48:30 charon: 15[ENC] parsed IKE_SA_INIT request 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) N(HASH_ALG) N(REDIR_SUP) ]
2016-06-15 12:48:30 charon: 15[IKE] received retransmit of request with ID 0, retransmitting response
2016-06-15 12:48:30 charon: 15[NET] sending packet: from fd07:7::ee[500] to fd07:7::e5[500] (584 bytes)
2016-06-15 12:48:37 charon: 14[NET] received packet: from fd07:7::e5[500] to fd07:7::ee[500] (1280 bytes)
2016-06-15 12:48:37 charon: 14[ENC] parsed IKE_SA_INIT request 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) N(HASH_ALG) N(REDIR_SUP) ]
2016-06-15 12:48:37 charon: 14[IKE] received retransmit of request with ID 0, retransmitting response
2016-06-15 12:48:37 charon: 14[NET] sending packet: from fd07:7::ee[500] to fd07:7::e5[500] (584 bytes)
2016-06-15 12:48:50 charon: 05[NET] received packet: from fd07:7::e5[500] to fd07:7::ee[500] (1280 bytes)
2016-06-15 12:48:50 charon: 05[ENC] parsed IKE_SA_INIT request 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) N(HASH_ALG) N(REDIR_SUP) ]
2016-06-15 12:48:50 charon: 05[IKE] received retransmit of request with ID 0, retransmitting response
2016-06-15 12:48:50 charon: 05[NET] sending packet: from fd07:7::ee[500] to fd07:7::e5[500] (584 bytes)
2016-06-15 12:48:56 charon: 09[JOB] deleting half open IKE_SA after timeout
2016-06-15 12:49:14 charon: 13[NET] received packet: from fd07:7::e5[500] to fd07:7::ee[500] (1280 bytes)
2016-06-15 12:49:14 charon: 13[ENC] parsed IKE_SA_INIT request 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) N(HASH_ALG) N(REDIR_SUP) ]
```

The initiator side:

```
conn ipsec2
 left=fd07:7::e5
 right=fd07:7::ee
 leftauth=psk
 rightauth=psk
 leftfirewall=yes
 keyexchange=ikev2
 ikelifetime=3600
 keylife=3600
 rekeymargin=540
 rekeyfuzz=100%
 keyingtries=%forever
 type=transport
 ike=aes128-sha256-modp3072,aes128-sha1-modp2048,3des-sha1-modp1536
 esp=aes128-sha256,aes128-sha1,3des-sha1
 auto=start

Status of IKE charon daemon (weakSwan 5.4.0, Linux 3.12.10+, armv7l):
  uptime: 16 minutes, since Jun 15 12:34:42 2016
  malloc: sbrk 405504, mmap 0, used 119208, free 286296
  worker threads: 11 of 16 idle, 5/0/0/0 working, job queue: 0/0/0/0, scheduled: 1
  loaded plugins: charon nonce pem openssl kernel-netlink socket-default stroke updown
Listening IP addresses:
  192.168.7.100
  fd07:7::e5
  192.168.2.1
  fd00:2::1
Connections:
      ipsec2:  fd07:7::e5...fd07:7::ee  IKEv2
      ipsec2:   local:  [fd07:7::e5] uses pre-shared key authentication
      ipsec2:   remote: [fd07:7::ee] uses pre-shared key authentication
      ipsec2:   child:  dynamic === dynamic TRANSPORT
Security Associations (0 up, 1 connecting):
      ipsec2[1]: CONNECTING, fd07:7::e5[%any]...fd07:7::ee[%any]
      ipsec2[1]: IKEv2 SPIs: fbff01842639b504_i* 0000000000000000_r
      ipsec2[1]: Tasks active: IKE_VENDOR IKE_INIT IKE_NATD IKE_CERT_PRE IKE_AUTH IKE_CERT_POST IK
E_CONFIG CHILD_CREATE IKE_AUTH_LIFETIME IKE_MOBIKE

2016-06-15 12:51:17 charon: 10[NET] sending packet: from fd07:7::e5[500] to fd07:7::ee[500] (1280
bytes)
2016-06-15 12:51:24 charon: 13[IKE] retransmit 2 of request with message ID 0
2016-06-15 12:51:24 charon: 13[NET] sending packet: from fd07:7::e5[500] to fd07:7::ee[500] (1280
bytes)
2016-06-15 12:51:37 charon: 06[IKE] retransmit 3 of request with message ID 0
2016-06-15 12:51:37 charon: 06[NET] sending packet: from fd07:7::e5[500] to fd07:7::ee[500] (1280
bytes)
```

This issue affects ikev1 as well.

History

#1 - 15.06.2016 15:30 - Tobias Brunner

- Description updated

- Category set to configuration

- Status changed from New to Feedback

> If I restart initiator side the connection is not automatically re-established.

Why should it? Or do you have a *conn %default* section that sets *closeaction=restart*?

```
2016-06-15 12:51:17 charon: 10[NET] sending packet: from fd07:7::e5[500] to fd07:7::ee[500] (1280 bytes)
2016-06-15 12:51:24 charon: 13[IKE] retransmit 2 of request with message ID 0

2016-06-15 12:48:30 charon: 15[NET] received packet: from fd07:7::e5[500] to fd07:7::ee[500] (1280 bytes)
2016-06-15 12:48:30 charon: 15[ENC] parsed IKE_SA_INIT request 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) N(HASH_A
LG) N(REDIR_SUP) ]
2016-06-15 12:48:30 charon: 15[IKE] received retransmit of request with ID 0, retransmitting response
2016-06-15 12:48:30 charon: 15[NET] sending packet: from fd07:7::ee[500] to fd07:7::e5[500] (584 bytes)
```

So it seems the IKE_SA_INIT response is dropped somewhere and does not reach the initiator. Try to find out where and why.

#2 - 16.06.2016 10:45 - Jiri Zendulka

Hi Tobias,

"Restart" means power cycle (switch off/on) in our case. Devices are routers. If I restart only ipsec service everything works fine. The responder got message about initiator restart so connection is deleted and initiator can initiate a new ipsec connection. But we concern about power cycle situation.

I found out a little bit different behaviour then the first one. The initiator sends neighbor solicitation to find a responder mac after power cycle. But the responder replies via original tunnel - it does not know that connection is not up anymore.

```
07:22:52.598994 IP6 fd07:7::e5 > ff02::1:ff00:ee: ICMP6, neighbor solicitation, who has fd07:7::ee, length 32
07:22:52.599606 IP6 fd07:7::ee > fd07:7::e5: ESP(spi=0xcb8de103,seq=0x76), length 88
```

Do you have any suggestion how to solve this issue?

IPv4 uses ARP which communicates on mac address level (not IP address). I guess that this is the reason why it works for ipv4 - the reply is not send via original tunnel.

#3 - 16.06.2016 11:09 - Tobias Brunner

> Do you have any suggestion how to solve this issue?

Yes, add passthrough policies for NDP messages:

```
conn ndp-ns
    right=::1 # so this connection does not get used for other purposes
    leftsubnet=::/0[ipv6-icmp/135]
    rightsubnet=::/0[ipv6-icmp/135]
    type=passthrough
    auto=route

conn ndp-na
    right=::1 # so this connection does not get used for other purposes
    leftsubnet=::/0[ipv6-icmp/136]
    rightsubnet=::/0[ipv6-icmp/136]
    type=passthrough
    auto=route
```

#4 - 16.06.2016 15:03 - Jiri Zendulka

Hi Tobias,

unfortunately it does not work for me. Are they any special rules in ip route or ip xfrm inserted by that ndp connections? Shall I check something?

```
Connections:
      ndp-ns:  ::1...%any  IKEv1/2
      ndp-ns:   local:  [::1] uses public key authentication
      ndp-ns:   remote: uses public key authentication
      ndp-ns:   child:   ===  PASS
      ndp-na:   child:   ===  PASS
      ipsec1:  fd07:7::ee...%any  IKEv1
      ipsec1:   local:  [fd07:7::ee] uses pre-shared key authentication
      ipsec1:   remote: uses pre-shared key authentication
      ipsec1:   child:  dynamic === dynamic TRANSPORT
Shunted Connections:
      ndp-ns:   ===  PASS
      ndp-na:   ===  PASS
Security Associations (1 up, 0 connecting):
      ipsec1[2]: ESTABLISHED 3 seconds ago, fd07:7::ee[fd07:7::ee]...fd07:7::e5[fd07:7::e5]
      ipsec1[2]: IKEv1 SPIs: 61eb62ca6866622c_i 8652924f165fa783_r*, pre-shared key reauthentication in 47 min
utes
      ipsec1[2]: IKE proposal: AES_CBC_128/HMAC_SHA2_256_128/PRF_HMAC_SHA2_256/MODP_3072
      ipsec1{2}:  INSTALLED, TRANSPORT, reqid 2, ESP SPIs: c662ebcc_i cd458fd2_o
      ipsec1{2}:  AES_CBC_128/HMAC_SHA2_256_128, 0 bytes_i, 0 bytes_o, rekeying in 44 minutes
      ipsec1{2}:   fd07:7::ee/128 === fd07:7::e5/128
```

#5 - 16.06.2016 17:20 - Tobias Brunner

```
    ...
          ndp-ns:   child:   ===  PASS
          ndp-na:   child:   ===  PASS
    ...
    Shunted Connections:
          ndp-ns:   ===  PASS
          ndp-na:   ===  PASS
```

Well, that looks wrong, doesn't it?

This is what it actually should look like (and does here when I test it):

```
...
      ndp-ns:   child:  ::/0[ipv6-icmp/135] === ::/0[ipv6-icmp/135] PASS
      ndp-na:   child:  ::/0[ipv6-icmp/136] === ::/0[ipv6-icmp/136] PASS
...
Shunted Connections:
      ndp-ns:  ::/0[ipv6-icmp/135] === ::/0[ipv6-icmp/135] PASS
      ndp-na:  ::/0[ipv6-icmp/136] === ::/0[ipv6-icmp/136] PASS
```

#6 - 16.06.2016 20:49 - Jiri Zendulka

Hi Tobias,

I managed it in the end. But I had to used proto number 58 instead of the proto name ipv6-icmp.

```
Shunted Connections:
      ndp-ns:  ::/0[58/135] === ::/0[58/135] PASS
      ndp-na:  ::/0[58/136] === ::/0[58/136] PASS
```

It works now.

Thanks.

#7 - 17.06.2016 09:10 - Tobias Brunner

- Status changed from Feedback to Closed

- Assignee set to Tobias Brunner

- Resolution set to No change required

> I managed it in the end. But I had to used proto number 58 instead of the proto name ipv6-icmp.

OK. We use getprotobyname() to parse that, could be that this doesn't work on your system or that this particular protocol is not defined (check /etc/protocols). And you probably got "invalid proto/port: ..., skipped subnet" error messages in the log while the config was loaded.