

## strongSwan - Bug #3511

### after IPv6 prefix change charon is stuck in "connecting"

12.07.2020 18:19 - Harald Dunkel

<b>Status:</b>	Closed	<b>Start date:</b>	
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Tobias Brunner	<b>Estimated time:</b>	0.00 hour
<b>Category:</b>	kernel-interface	<b>Resolution:</b>	Fixed
<b>Target version:</b>	5.9.1		
<b>Affected version:</b>	5.8.4		

**Description**

If I reboot my gateway to the internet, then Deutsche Telekom gives me a new IPv6 prefix. The old one is not routed anymore, even though its not expired yet. Can't help it.

Problem: charon tries to establish the connection to the IPsec gateway in the office using only the old prefix, even though both are known. ipsec statusall shows

```
# ipsec statusall
Status of IKE charon daemon (strongSwan 5.8.4, Linux 5.7.8-raw, x86_64):
  uptime: 4 hours, since Jul 12 13:58:38 2020
  malloc: sbrk 3219456, mmap 0, used 1330400, free 1889056
  worker threads: 11 of 16 idle, 5/0/0/0 working, job queue: 0/0/0/0, scheduled: 6
  loaded plugins: charon test-vectors ldap pkcs11 tpm aesni aes rc2 sha2 sha1 md5 mgf1 rdrand rand
om nonce x509 revocation constraints pubkey pkcs1 pkcs7 pkcs8 pkcs12 pgp dnskey sshkey pem openssl
  gcrypt af-alg fips-prf gmp curve25519 agent chapoly xcbc cmac hmac ctr ccm gcm drbg curl attr ker
nel-netlink resolve socket-default connmark farp stroke vici updown eap-identity eap-aka eap-md5 e
ap-gtc eap-radius eap-tls eap-ttls eap-tnc xauth-eap xauth-pam tnc-tncs dhcp lookip error-notify
  certexpire led addrblock unity counters
Listening IP addresses:
  10.42.100.186
  2003:e3:1f38:d803:ac5a:50ff:404b:2d4b
  2003:e3:1f40:6c03:766e:2cda:a51f:ac38
Connections:
cecil-ipsecgate: %any...ipsecgate.example.com IKEv2, dpddelay=30s
cecil-ipsecgate: local: [cecil.afaics.de] uses public key authentication
cecil-ipsecgate: cert: "C=DE, ST=NRW, L=Aachen, O=example AG, CN=cecil.afaics.de, E=security@e
xample.de"
cecil-ipsecgate: remote: [ipsecgate.example.com] uses public key authentication
cecil-ipsecgate: child: dynamic == 0.0.0.0/0 TUNNEL, dpdaction=restart
Security Associations (0 up, 1 connecting):
cecil-ipsecgate[3]: CONNECTING, 2003:e3:1f38:d803:ac5a:50ff:404b:2d4b[%any]...2001:db8:30:ffe0::d1
[%any]
cecil-ipsecgate[3]: IKEv2 SPIs: f3a281231d1d437c_i* 0000000000000000_r
cecil-ipsecgate[3]: Tasks active: IKE_VENDOR IKE_INIT IKE_NATD IKE_CERT_PRE IKE_AUTH IKE_CERT_POST
  IKE_CONFIG CHILD_CREATE IKE_AUTH_LIFETIME IKE_MOBIKE
```

Please note the old (2003:e3:1f38:d803::) and new (2003:e3:1f40:6c03::) prefix for the listen port.

Question is, shouldn't charon move to the new IPv6 address for outgoing connections asap, even though the old prefix is not expired yet?

#### Associated revisions

##### Revision 2eb43ca4 - 29.10.2020 09:46 - Tobias Brunner

kernel-netlink: Update cached address flags

Note that manually adding an IPv6 address without disabling duplicate address detection (DAD, e.g. via `nodad` when using iproute2) will cause a roam event due to a flag change after about 1-2 seconds (TENTATIVE is removed). If this is a problem, we might have to ignore addresses with TENTATIVE flag when we receive a RTM\_NEWADDR message until that flag is

eventually removed.

Fixes #3511.

### Revision a689e358 - 29.10.2020 09:46 - Tobias Brunner

kernel-netlink: Ignore deprecated candidate source addresses

The currently used address may get deprecated e.g. if an IPv6 prefix changes.  
In this case we should switch to another address.

Fixes #3511.

## History

---

### #1 - 13.07.2020 06:08 - Harald Dunkel

- File *charon.log.gz* added

### #2 - 13.07.2020 06:54 - Harald Dunkel

More information is necessary:

- This is a home office setup. The IPsec connection is established between my desktop PC and the gateway in the office, using the IPv6 prefix it has received via router advertisement from my gateway.
- Payload is IPv4 only.
- The problem comes up when charon tries to **reconnect** via IPv6.
- **ip a** shows the old IPv6 address as "deprecated". Its not expired, yet.

### #3 - 13.07.2020 07:03 - Harald Dunkel

PS: **ip a** shows this:

```
# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,PROMISC,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast master br0 state UP group default
    qlen 1000
    link/ether 0c:9d:92:66:99:ef brd ff:ff:ff:ff:ff:ff
    altnam eno2
    altnam enp0s31f6
3: wlan0: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
    link/ether 94:b8:6d:55:35:61 brd ff:ff:ff:ff:ff:ff
    altnam wlo1
    altnam wlp0s20f3
4: br0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1000
    link/ether 6a:fd:ff:6f:0e:76 brd ff:ff:ff:ff:ff:ff
    inet 10.42.100.186/24 brd 10.42.100.255 scope global noprefixroute br0
        valid_lft forever preferred_lft forever
    inet 172.19.122.225/32 scope global br0
        valid_lft forever preferred_lft forever
    inet6 2003:e3:1f40:6c03:766e:2cda:a51f:ac38/64 scope global dynamic mngtmpaddr noprefixroute
        valid_lft 3597sec preferred_lft 1797sec
    inet6 2003:e3:1f38:d803:ac5a:50ff:404b:2d4b/64 scope global deprecated dynamic mngtmpaddr noprefixroute
        valid_lft 3506sec preferred_lft 0sec
    inet6 fe80::4380:3a2a:7f76:29c4/64 scope link
        valid_lft forever preferred_lft forever
```

Please note the "preferred\_lft 0sec".

### #4 - 20.07.2020 16:25 - Tobias Brunner

- Category set to *kernel-interface*

- Status changed from *New* to *Feedback*

Question is, shouldn't charon move to the new IPv6 address for outgoing connections asap, even though the old prefix is not expired yet?

The default source address selection is mainly based on the routes, so if the old IP is still available and referenced in a route it has no reason to change anything (make sure that's not the case). Second, the current address is also tried to be kept (you can try enabling `charon.prefer_best_path` which makes strongSwan ignore any previous path/IP address if addresses change - but not when reestablishing/initiating and an address is already migrated from a previous SA). Afterwards, if multiple IPv6 source addresses are available (e.g. on the same interface to which a route points), deprecated addresses are generally avoided, however, currently not if we have a matching previous address that we can keep. I pushed a commit to the `3511-netlink-ignore-deprecated` branch that ignores such matches if the address is deprecated. Let me know if that helps.

Also, to maybe avoid all of the above, you could also try using the more efficient route/source lookup by setting `charon.plugins.kernel-netlink.fwmark` to `!<any numeric mark>`. Then the kernel handles these lookups and while we still pass the previous address for guidance it maybe ignores it if it's deprecated.

#### #5 - 22.07.2020 11:01 - Harald Dunkel

Thats a lot of ifs and maybes. Wouldn't you agree that the deprecated IPv6 address should be avoided automatically, if there is a better choice, no matter what?

Of course I applied the patch you provided. It builds fine, IPsec works, but I haven't had a chance to verify the prefix change yet, using the config options you suggested.

The patch alone seems to be insufficient (as expected). After rebooting the router it still got stuck in connecting. Using "ip a del" i removed the deprecated IPv6 address from the interface, but this did not help, even though the old IPv6 address doesn't show up in the routing table anymore.

#### #6 - 15.10.2020 16:15 - Tobias Brunner

I think I found the reason why the patch didn't help. The address flags were not actually updated when they changed, so the address was never marked deprecated in the internal address cache. I pushed another commit to the branch above that should address this.

#### #7 - 29.10.2020 09:47 - Tobias Brunner

- Tracker changed from Issue to Bug
- Assignee set to Tobias Brunner
- Target version set to 5.9.1
- Resolution set to Fixed

I did some tests with the [ipsec/rw-ikev2](#) scenario and adding an IPv6 address

```
carol# ip addr add fec0::42/16 dev eth0 nodad
```

```
carol charon-systemd: 09[KNL] fec0::42 appeared on eth0
carol charon-systemd: 05[IKE] sending address list update using MOBIKE
carol charon-systemd: 05[ENC] generating INFORMATIONAL request 2 [ N(ADD_4_ADDR) N(ADD_6_ADDR) ]
carol charon-systemd: 05[NET] sending packet: from fec0::10[4500] to fec0::1[4500] (112 bytes)
carol charon-systemd: 10[NET] received packet: from fec0::1[4500] to fec0::10[4500] (80 bytes)
carol charon-systemd: 10[ENC] parsed INFORMATIONAL response 2 [ ]
```

and deprecating the current one correctly switches to the new address:

```
carol# ip addr change fec0::10/16 dev eth0 preferred_lft 0
```

```
carol charon-systemd: 11[KNL] flags changed for fec0::10 on eth0
carol charon-systemd: 12[IKE] old path is not available anymore, try to find another
carol charon-systemd: 12[IKE] looking for a route to fec0::1 ...
carol charon-systemd: 12[IKE] requesting address change using MOBIKE
carol charon-systemd: 12[ENC] generating INFORMATIONAL request 3 [ ]
carol charon-systemd: 12[IKE] checking path fec0::42[4500] - fec0::1[4500]
carol charon-systemd: 12[NET] sending packet: from fec0::42[4500] to fec0::1[4500] (80 bytes)
carol charon-systemd: 12[IKE] checking path 192.168.0.100[4500] - 192.168.0.1[4500]
carol charon-systemd: 12[NET] sending packet: from 192.168.0.100[4500] to 192.168.0.1[4500] (80 bytes)
carol charon-systemd: 12[IKE] checking path 192.168.0.100[4500] - 10.1.0.1[4500]
carol charon-systemd: 12[NET] sending packet: from 192.168.0.100[4500] to 10.1.0.1[4500] (80 bytes)
carol charon-systemd: 16[NET] received packet: from fec0::1[4500] to fec0::42[4500] (80 bytes)
carol charon-systemd: 16[ENC] parsed INFORMATIONAL response 3 [ ]
carol charon-systemd: 16[ENC] generating INFORMATIONAL request 4 [ N(UPD_SA_ADDR) N(NATD_S_IP) N(NATD_D_IP) N(
```

```
COOKIE2) N(ADD_4_ADDR) ]
carol charon-systemd: 16[NET] sending packet: from fec0::42[4500] to fec0::1[4500] (176 bytes)
carol charon-systemd: 14[NET] received packet: from fec0::1[4500] to fec0::42[4500] (160 bytes)
carol charon-systemd: 14[ENC] parsed INFORMATIONAL response 4 [ N(NATD_S_IP) N(NATD_D_IP) N(COOKIE2) ]
```

So I'll include these changes in the next release.

**#8 - 29.10.2020 09:47 - Tobias Brunner**

- Status changed from Feedback to Closed

**#9 - 31.12.2020 21:32 - Harald Dunkel**

Sorry for not providing feedback. I forgot about it completely. Apparently the fix works. I just tried.

Thanx very much and best wishes for the new year. Please keep on your good work.

**#10 - 04.01.2021 12:08 - Tobias Brunner**

OK, great, thanks for testing.

**Files**

---

charon.log.gz	97 KB	13.07.2020	Harald Dunkel
---------------	-------	------------	---------------