strongSwan - Issue #2769

strongSwan is nearly unusable when the first dns server is down

23.09.2018 16:48 - Marcel Müller

Status:	Closed		
Priority:	Normal		
Assignee:	Tobias Brunner		
Category:	configuration		
Affected version:	5.6.2	**Resolution:**	No change required

Description

Hello everyone,

I've noticed on several occasions that my tunnels went down and were unable to get up again (strongSwan being the responder). After some testing I found out that this always happens when our main dns server went into maintenance and was unavailable. As the machine running strongSwan has 2 dns server configured in /etc/resolv.conf I never thought this would be a problem. I tried to compensate the issue with setting rotation and timeout options in resolv.conf, but this didn't help.

resolv.conf looks like this:

```
root@strongSwan:~# cat /etc/resolv.conf
domain <...>
search <...>
nameserver 172.31.1.9
nameserver 172.31.1.4
options timeout:1 rotate
```

All tunnels (IKEv1) are using a dyndns host as the rightid, like this:

```
conn 9022
        also=fritzbox
        right=9022.<domain>
        rightid="@9022.<domain>"
        rightsubnet=10.90.22.0/24
        auto=add
```

Now when the first nameserver is offline (172.31.1.9) more and more jobs are queued:

```
root@strongSwan:~# ipsec statusall
Status of IKE charon daemon (strongSwan 5.6.2, Linux 4.9.0-0.bpo.4-amd64, x86_64):
  uptime: 21 hours, since Sep 22 18:10:45 2018
  malloc: sbrk 8110080, mmap 0, used 4035888, free 4074192
  worker threads: 0 of 32 idle, 7/0/0/25 working, job queue: 0/61/6/750, scheduled: 426

root@strongSwan:~# ipsec statusall | grep worker
  worker threads: 0 of 32 idle, 8/0/0/24 working, job queue: 0/87/7/984, scheduled: 364

root@strongSwan:~# ipsec statusall | grep worker
  worker threads: 0 of 32 idle, 7/2/0/23 working, job queue: 0/12/1/2400, scheduled: 302
```

DNS resolution with nslookup still works as it uses the second nameserver:

```
root@strongSwan:~# nslookup google.de
Server:         172.31.1.4
Address:        172.31.1.4#53

Non-authoritative answer:
Name:   google.de
Address: 216.58.205.227
```

Processing of config candiates is very slow:

```
Sep 23 16:11:06 08[MGR] created IKE_SA (unnamed)[1388]
Sep 23 16:11:06 08[NET] <1388> received packet: from <ip>[500] to 172.31.1.5[500] (496 bytes)
Sep 23 16:11:06 08[ENC] <1388> parsed ID_PROT request 0 [ SA V V V V V V ]
Sep 23 16:11:06 08[CFG] <1388> looking for an ike config for 172.31.1.5...<ip>
Sep 23 16:11:07 30[CFG] <1343>   candidate: %any...9001.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:07 30[CFG] <1343>   candidate: %any...9005.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:08 30[CFG] <1343>   candidate: %any...9006.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:08 30[CFG] <1343>   candidate: %any...9009.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:09 30[CFG] <1343>   candidate: %any...9015.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:09 30[CFG] <1343>   candidate: %any...9018.<domain>,0.0.0.0/0,::/0, prio 2076
Sep 23 16:11:10 30[CFG] <1343>   candidate: %any...9019.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:10 30[CFG] <1343>   candidate: %any...9021.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:11 30[CFG] <1343>   candidate: %any...9022.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:11 30[CFG] <1343>   candidate: %any...9025.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:12 30[CFG] <1343>   candidate: %any...9031.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:12 30[CFG] <1343>   candidate: %any...9032.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:14 30[CFG] <1343>   candidate: %any...9036.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:14 30[CFG] <1343>   candidate: %any...9040.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:15 30[CFG] <1343>   candidate: %any...9041.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:15 30[CFG] <1343>   candidate: %any...9042.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:16 30[CFG] <1343>   candidate: %any...9045.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:16 30[CFG] <1343>   candidate: %any...9046.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:17 30[CFG] <1343>   candidate: %any...9049.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:18 30[CFG] <1343>   candidate: %any...9050.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:19 30[CFG] <1343>   candidate: %any...9051.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:19 30[CFG] <1343>   candidate: %any...9052.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:20 30[CFG] <1343>   candidate: %any...9053.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:20 30[CFG] <1343>   candidate: %any...9054.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:21 30[CFG] <1343>   candidate: %any...9055.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:21 30[CFG] <1343>   candidate: %any...9056.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:22 30[CFG] <1343>   candidate: %any...9057.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:22 30[CFG] <1343>   candidate: %any...9058.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:23 30[CFG] <1343>   candidate: %any...9059.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:23 30[CFG] <1343>   candidate: %any...9060.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:24 30[CFG] <1343>   candidate: %any...9102.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:24 30[CFG] <1343>   candidate: %any...9103.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:25 30[CFG] <1343>   candidate: %any...9107.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:25 30[CFG] <1343>   candidate: %any...9108.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:26 30[CFG] <1343>   candidate: %any...9201.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:27 30[CFG] <1343>   candidate: %any...9202.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:28 30[CFG] <1343>   candidate: %any...9203.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:28 30[CFG] <1343>   candidate: %any...9501.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:29 30[CFG] <1343>   candidate: %any...9503.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:29 30[CFG] <1343>   candidate: %any...9508.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:30 30[CFG] <1343>   candidate: %any...9509.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:30 30[CFG] <1343>   candidate: %any...9512.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:31 30[CFG] <1343>   candidate: %any...9514.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:31 30[CFG] <1343>   candidate: %any...9518.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:32 30[CFG] <1343>   candidate: %any...9520.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:32 30[CFG] <1343>   candidate: %any...9524.<domain>,0.0.0.0/0,::/0, prio 28
Sep 23 16:11:33 30[CFG] <1343>   candidate: %any...9525.<domain>,0.0.0.0/0,::/0, prio 28
```

ipsec statusall completed after 14 minutes (and only after the main nameserver was up again).

strongswan.conf looks like this:

```
charon {

        # number of worker threads in charon
        threads = 32
        host_resolver.max_threads = 10

(....)
```

```
}
```

under normal circumstances the worker threads/queue looks like this

```
root@strongSwan:~# date && ipsec statusall | grep worker
So 23. Sep 16:24:41 CEST 2018
  worker threads: 27 of 32 idle, 5/0/0/0 working, job queue: 0/0/0/0, scheduled: 476
```

To be fair, I'm not sure if this is a strongSwan issue or if this "works as designed". But I don't think the right way to compensate this issue is to increase the number of worker threads, is it? Can the selection of which nameserver to use be optimized? Is it strongSwans choice at all?

Thanks in advance,
Marcel

History

#1 - 24.09.2018 10:51 - Tobias Brunner

- Status changed from New to Feedback

> Processing of config candiates is very slow:

Because this requires at least one DNS resolution for each candidate if FQDNs are configured as local/remote addresses (*rightid* is not relevant, *right* is) to find the best match (based on the IP addresses of the received packet and the identities, which are never resolved).

> Can the selection of which nameserver to use be optimized? Is it strongSwans choice at all?

No, strongSwan uses the getaddrinfo(3) function to resolve hostnames. So it's up to libc's resolver which DNS servers are used. Using a different local resolver (via libc resolver directed to 127.0.0.1) like dnsmasq or unbound might help (e.g. in regards to switching between servers or caching the results).

#2 - 24.09.2018 14:53 - Marcel Müller

Hello Tobias,

thanks for your quick response!
I was under the impression that when using DynDNS FQDNs as remote addresses that --auto-update 300 needs to be used in the start script as otherwise the FQDNs would only be resolved at the start of charon? Therefore I was wondering why a DNS query is needed for each config at each connect (50 tunnels, 50 configs -> 2500 queries?). Maybe I'm mixing things up... Sorry if that's the case.

I like your idea of a local dns resolver, thanks for that. I think I'll go this route.

#3 - 24.09.2018 15:31 - Tobias Brunner

> I was under the impression that when using DynDNS FQDNs as remote addresses that --auto-update 300 needs to be used in the start script as otherwise the FQDNs would only be resolved at the start of charon?

I don't think this was ever necessary with charon (certainly not since 5.0.0). It pretty much always did the DNS resolution itself and not when loading a config. The old IKEv1 daemon pluto, however, probably did require this. The starter process, which (re-)loads the ipsec.conf file, also does no DNS resolution anymore (was removed with 5.0.0), so if nothing else changed in the config file, the auto update is basically a no-op.

> (50 tunnels, 50 configs -> 2500 queries?)

Yes, basically. However, it depends on the caching whether each requires actually querying the DNS server (strongSwan does no caching itself, though, so that's up to the resolver).

> I like your idea of a local dns resolver, thanks for that. I think I'll go this route.

Please report back if this works out, or with possible caveats.

#4 - 08.10.2018 17:39 - Marcel Müller

Hello Tobias,

thanks for the info about --auto-update!
I've been running strongswan with unbound installed locally for about 2 weeks now and it works great. My config looks like this:

```
root@strongSwan:~# cat /etc/resolv.conf
domain <myDomain>
search <myDomain>
nameserver 127.0.0.1

root@strongSwan:~# cat /etc/unbound/unbound.conf
include: "/etc/unbound/unbound.conf.d/*.conf"

server:
        hide-identity: yes
        hide-version: yes
        access-control: 127.0.0.1/32 allow
        cache-min-ttl: 60
        minimal-responses: yes
        num-threads: 4
        prefetch: yes
        domain-insecure: "<myDomain>"

forward-zone:
        name: "<myDomain>."
        forward-addr: 172.31.1.9
        forward-addr: 172.31.1.4
        forward-first: yes

forward-zone:
        name: "."
        forward-addr: 1.1.1.1        # Cloudflare
        forward-addr: 1.0.0.1        # Cloudflare
        forward-addr: 8.8.4.4        # Google
        forward-addr: 8.8.8.8        # Google
```

Unbound allows to set a minimum for ttls which can be used to increase the cache time for dyndns hosts.

In the 2 weeks our main dns was shut down for maintenance and strongswan kept working perfectly.

Thanks again,
Marcel

#5 - 08.10.2018 17:47 - Tobias Brunner

- Category set to configuration

- Status changed from Feedback to Closed

- Assignee set to Tobias Brunner

- Resolution set to No change required

Great, thanks for the update.