strongSwan - Issue #1501

IPsec tunnel is not automatically reestablished if the remote side is restarted

07.06.2016 11:39 - Jiri Zendulka

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

Description

Hello,

I established an ikev1 ipsec connection beetwen two strongswans. One side is configured as initiator and the second one is configured as responder.

The initiator ipsec.conf:

```
conn ipsec1
 left=192.168.7.110
 right=192.168.7.120
 leftauth=psk
 rightauth=psk
 leftsubnet=192.168.100.0/24
 rightsubnet=192.168.1.0/24
 leftfirewall=yes
 keyexchange=ikev1
 ikelifetime=3600
 keylife=3600
 rekeymargin=540
 rekeyfuzz=100%
 keyingtries=%forever
 type=tunnel
 auto=start
 closeaction=restart
```

The responder ispec.conf:

```
conn ipsec1
 left=192.168.7.120
 right=%any
 leftauth=psk
 rightauth=psk
 leftsubnet=192.168.1.0/24
 rightsubnet=192.168.100.0/24
 leftfirewall=yes
 keyexchange=ikev1
 ikelifetime=3600
 keylife=3600
 rekeymargin=540
 rekeyfuzz=100%
 keyingtries=%forever
 type=tunnel
 auto=add
```

I found the following issue: IPsec tunnel is not automatically reestablished if the remote side is restarted, so manual action on the local side is required. Option closeaction=restart should be used in ipsec.conf in this case, but it is not recommended to use it with uniqueids=yes that is usually enabled. Unfortunatelly even closeaction=restart does not help if tunnel is downed and deleted on the remote side, because the remote side sends NO_PROPOSAL_CHOSEN and SA is permanently deleted on the local side too.

Many thanks for your help.

History

#1 - 07.06.2016 15:32 - Tobias Brunner

- Description updated

- Category set to configuration

- Status changed from New to Feedback

> because the remote side sends NO_PROPOSAL_CHOSEN

Why is that? Is the config not loaded yet?

Also if you use strongSwan on both sides, please use IKEv2. And *auto=route* might work better if you want the connection to be always up.

#2 - 08.06.2016 08:02 - Jiri Zendulka

Hi Tobias,

We need to use ikev1 too. Some customers still use ikev1. The responder side is stopped by commands stroke down + stroke delete and then the responder side is started up by starter cmd. The initiator is not able re-established connection after that. The responder sends NO_PROPOSAL_CHOSEN and SA is deleted at initiator side then.

Thanks.

#3 - 08.06.2016 09:50 - Tobias Brunner

> The responder side is stopped by commands stroke down + stroke delete and then the responder side is started up by starter cmd. The initiator is not able re-established connection after that. The responder sends NO_PROPOSAL_CHOSEN and SA is deleted at initiator side then.

Why does it send that notify? Is the connection loaded again? Why do you use stroke directly in the first place?

#4 - 08.06.2016 10:24 - Jiri Zendulka

...stroke can't be used directly? Is something wrong to terminate the connection via stroke down + stroke delete? We use strongswan at a linux based embedded device and we need to save flash memory size as much as possible. So the ipsec script is not installed currently.

The ipsec responder status after restart:

```
Status of IKE charon daemon (weakSwan 5.4.0, Linux 3.5.0-lsp-3.3.1, armv5tejl):
  uptime: 24 seconds, since Jun 08 10:11:14 2016
  malloc: sbrk 405504, mmap 0, used 106752, free 298752
  worker threads: 11 of 16 idle, 5/0/0/0 working, job queue: 0/0/0/0, scheduled: 0
  loaded plugins: charon nonce pem openssl kernel-netlink socket-default stroke updown
Listening IP addresses:
  192.168.1.1
  192.168.7.120
  192.168.3.1
Connections:
      ipsec1:  192.168.7.120...%any  IKEv1
      ipsec1:   local:  [192.168.7.120] uses pre-shared key authentication
      ipsec1:   remote: uses pre-shared key authentication
      ipsec1:   child:  192.168.1.0/24 === 192.168.100.0/24 TUNNEL
Security Associations (0 up, 0 connecting):
  none
```

The responder log:

```
....
2016-06-08 10:11:14 charon: 07[CFG] received stroke: add connection 'ipsec1'
2016-06-08 10:11:14 charon: 07[CFG] conn ipsec1
2016-06-08 10:11:14 charon: 07[CFG]   left=192.168.7.120
2016-06-08 10:11:14 charon: 07[CFG]   leftsubnet=192.168.1.0/24
2016-06-08 10:11:14 charon: 07[CFG]   leftauth=psk
2016-06-08 10:11:14 charon: 07[CFG]   leftupdown=/etc/scripts/updown
2016-06-08 10:11:14 charon: 07[CFG]   right=%any
2016-06-08 10:11:14 charon: 07[CFG]   rightsubnet=192.168.100.0/24
2016-06-08 10:11:14 charon: 07[CFG]   rightauth=psk
```

```
2016-06-08 10:11:14 charon: 07[CFG]   ike=aes128-sha256-modp3072
2016-06-08 10:11:14 charon: 07[CFG]   esp=aes128-sha256
2016-06-08 10:11:14 charon: 16[LIB] created thread 16 [1078]
2016-06-08 10:11:14 charon: 16[JOB] started worker thread 16
2016-06-08 10:11:14 charon: 15[LIB] created thread 15 [1079]
2016-06-08 10:11:14 charon: 15[JOB] started worker thread 15
2016-06-08 10:11:14 charon: 03[JOB] watcher going to poll() 4 fds
2016-06-08 10:11:14 charon: 07[CFG]   dpddelay=30
2016-06-08 10:11:14 charon: 07[CFG]   dpdtimeout=150
2016-06-08 10:11:14 charon: 07[CFG]   mediation=no
2016-06-08 10:11:14 charon: 07[CFG]   keyexchange=ikev1
2016-06-08 10:11:14 charon: 07[CFG] added configuration 'ipsec1'
2016-06-08 10:11:14 charon: 03[JOB] watcher got notification, rebuilding
2016-06-08 10:11:14 charon: 03[JOB] watcher going to poll() 4 fds
2016-06-08 10:11:17 charon: 03[JOB] watched FD 12 ready to read
2016-06-08 10:11:17 charon: 03[JOB] watcher going to poll() 3 fds
2016-06-08 10:11:17 charon: 03[JOB] watcher got notification, rebuilding
2016-06-08 10:11:17 charon: 03[JOB] watcher going to poll() 4 fds
2016-06-08 10:11:17 charon: 08[CFG] proposing traffic selectors for us:
2016-06-08 10:11:17 charon: 08[CFG]  192.168.1.0/24
2016-06-08 10:11:17 charon: 08[CFG] proposing traffic selectors for other:
2016-06-08 10:11:17 charon: 08[CFG]  192.168.100.0/24
2016-06-08 10:11:17 charon: 03[JOB] watcher got notification, rebuilding
2016-06-08 10:11:17 charon: 03[JOB] watcher going to poll() 4 fds
2016-06-08 10:11:18 charon: 03[JOB] watched FD 12 ready to read
2016-06-08 10:11:18 charon: 03[JOB] watcher going to poll() 3 fds
2016-06-08 10:11:18 charon: 03[JOB] watcher got notification, rebuilding
2016-06-08 10:11:18 charon: 03[JOB] watcher going to poll() 4 fds
2016-06-08 10:11:18 charon: 09[CFG] proposing traffic selectors for us:
2016-06-08 10:11:18 charon: 09[CFG]  192.168.1.0/24
2016-06-08 10:11:18 charon: 09[CFG] proposing traffic selectors for other:
2016-06-08 10:11:18 charon: 09[CFG]  192.168.100.0/24
2016-06-08 10:11:18 charon: 03[JOB] watcher got notification, rebuilding
....
```

So I think that configuration is properly loaded after responder restarting.

The ipsec initiator status after responder restart:

```
Status of IKE charon daemon (weakSwan 5.4.0, Linux 3.5.0-lsp-3.3.1, armv5tejl):
  uptime: 13 minutes, since Jun 08 08:11:16 2016
  malloc: sbrk 405504, mmap 0, used 106480, free 299024
  worker threads: 11 of 16 idle, 5/0/0/0 working, job queue: 0/0/0/0, scheduled: 2
  loaded plugins: charon nonce pem openssl kernel-netlink socket-default stroke updown
Listening IP addresses:
  192.168.100.1
  192.168.7.110
Connections:
      ipsec1:  192.168.7.110...192.168.7.120  IKEv1
      ipsec1:   local:  [192.168.7.110] uses pre-shared key authentication
      ipsec1:   remote: [192.168.7.120] uses pre-shared key authentication
      ipsec1:   child:  192.168.100.0/24 === 192.168.1.0/24 TUNNEL
Security Associations (0 up, 0 connecting):
  none
```

The initiator log:

```
2016-06-08 08:11:20 charon: 14[ENC] generating QUICK_MODE request 3414626679 [ HASH ]
2016-06-08 08:11:20 charon: 14[NET] sending packet: from 192.168.7.110[500] to 192.168.7.120[500] (76 bytes)
2016-06-08 08:14:52 charon: 07[NET] received packet: from 192.168.7.120[500] to 192.168.7.110[500] (92 bytes)
2016-06-08 08:14:52 charon: 07[ENC] parsed INFORMATIONAL_V1 request 1226969262 [ HASH D ]
2016-06-08 08:14:52 charon: 07[IKE] received DELETE for ESP CHILD_SA with SPI c52f7024
2016-06-08 08:14:52 charon: 07[IKE] closing CHILD_SA ipsec1{1} with SPIs cf07063b_i (0 bytes) c52f7024_o (0 by
tes) and TS 192.168.100.0/24 === 192.168.1.0/24
2016-06-08 08:14:52 charon: 07[ENC] generating QUICK_MODE request 427754279 [ HASH SA No ID ID ]
2016-06-08 08:14:52 charon: 07[NET] sending packet: from 192.168.7.110[500] to 192.168.7.120[500] (220 bytes)
2016-06-08 08:14:52 charon: 11[NET] received packet: from 192.168.7.120[500] to 192.168.7.110[500] (108 bytes)

2016-06-08 08:14:52 charon: 11[ENC] parsed INFORMATIONAL_V1 request 2641507477 [ HASH D ]
```

"A moment when responder has been restarted"

```
2016-06-08 08:14:52 charon: 11[IKE] received DELETE for IKE_SA ipsec1[1]
2016-06-08 08:14:52 charon: 11[IKE] deleting IKE_SA ipsec1[1] between 192.168.7.110[192.168.7.110]...192.168.7
.120[192.168.7.120]
```

```
2016-06-08 08:14:52 charon: 11[IKE] initiating Main Mode IKE_SA ipsec1[2] to 192.168.7.120
2016-06-08 08:14:52 charon: 11[ENC] generating ID_PROT request 0 [ SA V V V ]
2016-06-08 08:14:52 charon: 11[NET] sending packet: from 192.168.7.110[500] to 192.168.7.120[500] (204 bytes)
2016-06-08 08:14:53 charon: 13[NET] received packet: from 192.168.7.120[500] to 192.168.7.110[500] (40 bytes)
2016-06-08 08:14:53 charon: 13[ENC] parsed INFORMATIONAL_V1 request 331690821 [ N(NO_PROP) ]
2016-06-08 08:14:53 charon: 13[IKE] received NO_PROPOSAL_CHOSEN error notify
```

#5 - 08.06.2016 10:47 - Tobias Brunner

> ...stroke can't be used directly?

You shouldn't. See stroke.

> Is something wrong to terminate the connection via stroke down + stroke delete?

Yes, because the starter daemon won't know the connection is deleted. stroke delete is never called by the ipsec script, connections are managed by the starter daemon.

> We use strongswan at a linux based embedded device and we need to save flash memory size as much as possible. So the ipsec script is not installed currently.

Wow, you saved about 7 KB.

You might want to consider using vici/swanctl instead of ipsec/starter/stroke.

> So I think that configuration is properly loaded after responder restarting.

But does the Main Mode request from the initiator perhaps arrive before it is?

#6 - 10.06.2016 15:30 - Jiri Zendulka

Hi Tobias,

I do not use "stroke delete" anymore as you wrote above and tunnel is succesfully re-established if the responder side is restarted.

But I found next issue for my application. I have a situation where tunnel beetween responder and initiator side is up. If I change i.e. ESP settings at responder side and then I restart the responder side the tunnel is not re-established - it is supposed = ESP misconfiguration. But when I set the right ESP settings again the initiator is not trying to connect anymore. When I use openswan in this situation the tunnel is successfully re-established.

I generally mean the situation when strongswan's initiators are set to one responder side (i.e.cisco) and someone changes ipsec configuration at the responder side and then changes to the correct one. Strongswan's initiators do not re-established tunnel to the cisco anymore.

If I am right the suggested auto=route does not solve this issue generally. For example: Passive PLC devices (respond to a server request only) behind strongswan's initiator (our embedded device) waiting for a request from server which is (ipsec) responder side. The request never comes to the PLC due to tunnel is down.

Thanks.

#7 - 10.06.2016 17:02 - Tobias Brunner

> When I use openswan in this situation the tunnel is successfully re-established.

That's not how strongSwan behaves if there is a fatal error like NO_PROPOSAL_CHOSEN. These are not expected to get fixed without manual intervention.

> I generally mean the situation when strongswan's initiators are set to one responder side (i.e.cisco) and someone changes ipsec configuration at the responder side and then changes to the correct one.

Why should that happen? If that's a valid concern you should consider who has access to your devices. Nobody that doesn't understand the implications should be allowed to change such configurations.

> If I am right the suggested auto=route does not solve this issue generally.

Well, it prevents plaintext traffic from leaving the host and it will trigger another acquire and a retry after a while (see

charon.plugins.kernel-netlink.xfrm_acq_expires in strongswan.conf).

#8 - 12.06.2016 20:57 - Jiri Zendulka

Hi Tobias,

we have been testing an strongswan's ability to deal with non-standard situation and comparing strongswan's behaviour with openswan. It is very important for us to not give up trying to connect to responder side due to temporary server (responder) misconfiguration. Manual intervetion is sometimes very problematic (many remote devices). As I mention above openswan is able to deal with that situation. We use openswan currently but we are going to replace it with strongswan due to better ikev2 implementation etc.

Is any possibility how to set strongswan initiator to not give up trying to connect to responder side in that situation?

Thanks.

#9 - 13.06.2016 09:12 - Tobias Brunner

> Is any possibility how to set strongswan initiator to not give up trying to connect to responder side in that situation?

No, currently not.

#10 - 13.06.2016 09:35 - Jiri Zendulka

...and we can expect that strongswan will have this option in near future? Or do you think that it is not usefull option?

#11 - 13.06.2016 09:53 - Tobias Brunner

> ...and we can expect that strongswan will have this option in near future? Or do you think that it is not usefull option?

I currently don't see us doing that. Several *libcharon* users rely on the fact that fatal failures (like NO_PROPOSAL_CHOSEN, authentication failures) don't trigger a retry.

But as mentioned, changing the configuration (*auto=route*) might mitigate the issue for you. You could perhaps also use VICI to write a wrapper when initiating the connection.

#12 - 13.06.2016 10:25 - Jiri Zendulka

I think that changing the configuration to auto=route does not solve following situation:
Passive PLC devices (responds to a server request only) behind strongswan's initiator (our embedded device) waiting for a request from server which is (ipsec) responder side. The request from server never comes to the PLC due to tunnel is down.
Am I right?

Could you give an advice how to patch libcharon to trigger a retry in case of this failures?

Thanks.

#13 - 13.06.2016 10:55 - Tobias Brunner

> I think that changing the configuration to auto=route does not solve following situation:
> Passive PLC devices (responds to a server request only) behind strongswan's initiator (our embedded device) waiting for a request from server which is (ipsec) responder side. The request from server never comes to the PLC due to tunnel is down.
> Am I right?

Sure, but you could always create traffic yourself to the responder (e.g. regular ICMP packets). But why not make the other devices the initiators? And why is that really an issue for you? Do you actually expect config changes to mess with your configuration? Why?

> Could you give an advice how to patch libcharon to trigger a retry in case of this failures?

Since there are numerous possible fatal errors there will be several places and recovery might be different in different cases. Have a look at the tasks in source:src/libcharon/sa/ikev2/tasks. Depending on what you want to do you could perhaps also use the hooks on bus_t and reinitiate the connection (or as mentioned use VICI/swanctl to reinitiate if initiation failed).

#14 - 13.06.2016 13:01 - Jiri Zendulka

There are lots of ipsec configurations how customers ipsec use. Any manual intervention (ipsec restart or create some traffic by icmp)is not possible

in real situation because our device are remote wireless routers:

Customer's PLC devices --- eth --- Wireless routers with dynamic ip (initiators) ---- mobile connetcion --- Customer's cisco server (responder) with public IP.

We are not administrators of that cisco server. We only produce the wireless routers and one of our customers uses this solution.
We do not suppose that someone will regularly change the ipsec configuration on the server side. But as known the world this things happens from time to time (in error).

We've just finish ipsec configuration via starter/stroke interface. Using VIC/swanctl instead of starter/stroke means strat again from beginning.

#15 - 13.06.2016 14:11 - Tobias Brunner

> We do not suppose that someone will regularly change the ipsec configuration on the server side. But as known the world this things happens from time to time (in error).

In which case you could also let them reboot the wireless router after they fixed the config.

> We've just finish ipsec configuration via starter/stroke interface. Using VIC/swanctl instead of starter/stroke means strat again from beginning.

You can also initiate connections defined in ipsec.conf via swanctl. But you should consider switching as starter/stroke will disappear in the long run.

#16 - 13.06.2016 15:18 - Jiri Zendulka

Rebooting router is the issue at that moment because manual intervention is needed. When the ipsec connection is not up. There is not remote access to the router. The router has a dynamic ip. And on top of that there are lots of routers connected to the server. Not only one.

We thing that strongswan should keep trying to connect to the responder even if receives NO_PROPOSAL_CHOSEN. Openswan works so.

Thanks

#17 - 13.06.2016 15:20 - Tobias Brunner

> We thing that strongswan should keep trying to connect to the responder even if receives NO_PROPOSAL_CHOSEN.

There are currently no plans to change the current behavior.

> Openswan works so.

Then use Openswan.

#18 - 15.06.2016 10:34 - Jiri Zendulka

- File strongswan-5.4.0-auto-restart.patch added

Hi Tobias,

we made a patch for checking/restarting connection if the connection is not up. What's your opinion of this patch? We think that this behaviour is usefull if ipsec is used on remote devices which are connected to mobile network (dynamic ip) and access to devices is dependent on ipsec connection.

Thanks.

#19 - 15.06.2016 14:51 - Tobias Brunner

> What's your opinion of this patch?

You could probably do the same without a patch using a script that calls ipsec status and ipsec up (or as mentioned before via swanctl/vici). But do whatever works for you.

#20 - 15.06.2016 15:07 - Jiri Zendulka

Thanks. You can close the issue now.

#21 - 15.06.2016 15:21 - Tobias Brunner

- Status changed from Feedback to Closed

- Assignee set to Tobias Brunner

- Resolution set to No change required

Files

strongswan-5.4.0-auto-restart.patch	2.87 KB	15.06.2016	Jiri Zendulka