pfSense - Bug #10175

VTI tunnels to AWS drop and do not automatically reconnect

01/10/2020 03:47 AM - Brian Candler

Status:	Duplicate	**Start date:**	01/10/2020
Priority:	Normal	**Due date:**	
Assignee:		**% Done:**	0%
Category:	IPsec	**Estimated time:**	0.00 hour
Target version:			
Affected Version:	2.4.4-p3	**Affected Architecture:**	

Description

On a HA pair of XG-1537, I have four VTI tunnels to AWS - two each to two different accounts, with BGP failover on each pair (OpenBGPD). I am also monitoring them with NRPE, using check_ping to the 169.254.x.x remote endpoint addresses.

Every few days, one of the tunnels stops working: check_ping reports the connection is down. If I don't fix this, then within a few hours AWS also sends me an E-mail alert telling me about loss of redundancy.

If I go into the IPSEC status page, I see the tunnel status as "Disconnected" with a green "Connect" button next to it. Clicking the button fixes the problem.

I have had a look through the ipsec_status.php code to see what it does, to replicate it at the command line. I've found that:

(1) When the tunnel is down, /usr/local/sbin/swanctl --list-sas does not show the affected SA at all (no entry for conX000)
(2) I can bring the tunnel back up by running /usr/local/sbin/swanctl --initiate --child conX000

Therefore, I now have a workaround in the form of a cronjob in /etc/cron.d/

```
0 * * * * root for c in con4000 con5000 con6000 con7000; do /usr/local/sbin/swanctl --list-sas | grep "$c" >/dev/null || /usr/local/sbin/swanctl --initiate --child "$c"; done
```

However if I can help to fix the underlying problem, I would like to do so. For example, if you give me any debugging commands you want me to run when the tunnel next fails, I can disable my workaround.

Looking through existing issues, #9767 may be related. Note that both OpenBGPD and NRPE should be generating background traffic all the time, even on the tunnel which is not carrying traffic.

Here is example --list-sas output for one tunnel when it is up:

```
con4000: #4366, ESTABLISHED, IKEv1, XXXXXXXX_i* XXXXXXXX_r
  local  'X.X.X.X' @ X.X.X.X[4500]
  remote '34.251.125.152' @ 34.251.125.152[4500]
  AES_CBC-128/HMAC_SHA1_96/PRF_HMAC_SHA1/MODP_1024
  established 5354s ago, reauth in 22381s
  con4000: #306875, reqid 4000, REKEYED, TUNNEL-in-UDP, ESP:AES_CBC-128/HMAC_SHA1_96/MODP_1024
    installed 2830s ago, rekeying in -309s, expires in 770s
    in  c685e90c,  34243 bytes,   515 packets
    out 737beb3c,  77536 bytes,   590 packets
    local  0.0.0.0/0|/0
    remote 0.0.0.0/0|/0
  con4000: #306964, reqid 4000, INSTALLED, TUNNEL-in-UDP, ESP:AES_CBC-128/HMAC_SHA1_96/MODP_1024
    installed 308s ago, rekeying in 2211s, expires in 3292s
    in  c4a46c3b,   4166 bytes,    63 packets
    out 47fb2cc0,   9200 bytes,    70 packets
    local  0.0.0.0/0|/0
    remote 0.0.0.0/0|/0
```

Here is the pfSense configuration for this tunnel (phase1, phase2 and NRPE):

```
        <phase1>
            <ikeid>4</ikeid>
            <iketype>ikev1</iketype>
            <mode>main</mode>
            <interface>_vip5ce58f3a60ba7</interface>
            <remote-gateway>34.251.125.152</remote-gateway>
            <protocol>inet</protocol>
            <myid_type>myaddress</myid_type>
            <myid_data></myid_data>
            <peerid_type>peeraddress</peerid_type>
            <peerid_data></peerid_data>
            <encryption>
                <item>
                    <encryption-algorithm>
                        <name>aes</name>
                        <keylen>128</keylen>
                    </encryption-algorithm>
                    <hash-algorithm>sha1</hash-algorithm>
                    <dhgroup>2</dhgroup>
                </item>
            </encryption>
            <lifetime>28800</lifetime>
            <pre-shared-key>XXXXXXXX</pre>
            <private-key></private-key>
            <certref></certref>
            <caref></caref>
            <authentication_method>pre_shared_key</authentication_method>
            <descr><![CDATA[AWS VPN 1]]></descr>
            <nat_traversal>on</nat_traversal>
            <mobike>off</mobike>
            <margintime>600</margintime>
            <dpd_delay>10</dpd_delay>
            <dpd_maxfail>3</dpd_maxfail>
        </phase1>
...
        <phase2>
            <ikeid>4</ikeid>
            <uniqid>5d2468911ce1d</uniqid>
            <mode>vti</mode>
            <reqid>2</reqid>
            <localid>
                <type>network</type>
                <address>169.254.23.10</address>
                <netbits>30</netbits>
            </localid>
            <remoteid>
                <type>address</type>
                <address>169.254.23.9</address>
            </remoteid>
            <protocol>esp</protocol>
            <encryption-algorithm-option>
                <name>aes</name>
                <keylen>128</keylen>
            </encryption-algorithm-option>
            <hash-algorithm-option>hmac_sha1</hash-algorithm-option>
            <pfsgroup>2</pfsgroup>
            <lifetime>3600</lifetime>
            <pinghost></pinghost>
            <descr></descr>
        </phase2>
...
                <row>
                    <name>check_aws_vpn1</name>
                    <command>check_ping</command>
                    <warning>500,50%</warning>
                    <critical>1000,100%</critical>
                    <extra>-H 169.254.23.9 -p 2 -t 3</extra>
                </row>
```

History

#1 - 01/10/2020 05:08 AM - Brian Candler

Note that in the above, the closing </pre-shared-key> tag was mangled by redmine to just </pre>

#2 - 01/10/2020 06:28 AM - Jim Pingle

- Category set to IPsec

- Status changed from New to Duplicate

Duplicate of #9767

Please post on the forum to discuss issues before opening bug reports, and search for existing issues before opening new ones.