

strongSwan - Issue #3291

IPSec IKEv2 Client to VPN service 2

12.12.2019 05:58 - Bernd Bernikov

Status: Feedback	
Priority: Normal	
Assignee: Tobias Brunner	
Category: configuration	
Affected version: 5.6.3	Resolution:
Description	
Tobias Brunner wrote:	
You need to provide more information for us to help (see HelpRequests).	
I try to describe my problem in more detail.	
The goal is an IKEv2 connection with an IPv6 IP and KillSwitch. (KillSwitch: Internet access should only work with IKEv2 connection.)	
From the VPN provider I have this Linux configuration:	
<pre>conn PP keyexchange=ikev2 dpdaction=none dpddelay=300s inactivity=36000s rekey=no leftsourceip=%config4,%config6 leftsendcert=never leftauth=eap-mschapv2 rightauth=pubkey right=amsterdam.perfect-privacy.com rightid=%any rightca=/etc/ipsec.d/cacerts/perfect-privacy_ipsec_ca.crt rightsubnet=0.0.0.0/0,::/0 rightsendcert=always eap_identity="Username" type=tunnel auto=add</pre>	
This configuration does not work with OpenWRT, so i modified this configuration from this thread: https://forum.openwrt.org/t/strongswan-client-router-no-lan-access-when-vpn-connects/44350	
Now I have this configuration:	
/etc/ipsec.conf:	
<pre>config setup # strictcrlpolicy=yes uniqueids=never # yes #uniqueids=never charondebug="all" #charondebug="cfg 3, dmn 4, ike 3, net 1" # Add connections here. conn lan-passthrough leftsubnet=192.168.1.0/24 # Replace with your LAN subnet rightsubnet=192.168.1.0/24 # Replace with your LAN subnet authby=never # No authentication necessary type=pass # passthrough auto=route # no need to ipsec up lan-passthrough</pre>	

```
conn PP
  eap_identity="Username"
  type=tunnel
  keyexchange=ikev2
  dpdaction=restart
  closeaction=restart
  dpddelay=300s
  inactivity=36000s
  rekey=no
  forceencaps=yes
  authby=secret
  ike=aes256-sha256-modp2048
  esp=aes256-sha256
  leftfirewall=yes
  left=192.168.1.1
  leftid=192.168.1.1
  leftsourceip=%config4,%config6
  leftsendcert=never
  leftauth=eap-mschapv2
  rightfirewall=yes
  rightauth=pubkey
  right=37.48.94.1
  rightid=%any
  rightsubnet=0.0.0.0/0,::/0
  rightsendcert=always
  auto=add
```

/etc/ipsec.secrets:

```
Username : EAP "Password"
```

/etc/ipsec.user:

```
case "$PLUTO_VERB" in
up-client)
  iptables -t nat -A postrouting_wan_rule -s 192.168.1.0/24 -m policy --dir out --pol none -
j SNAT --to-source "$PLUTO_MY_SOURCEIP4_1"
  ;;
down-client)
  iptables -t nat -F postrouting_wan_rule
  ;;
esac
```

root@OpenWrt:~# ipsec up PP

```
initiating IKE_SA PP[1] to 37.48.94.1
generating IKE_SA_INIT request 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) N(FRAG_SUP) N(HASH_ALG) N(RE
DIR_SUP) ]
sending packet: from 192.168.1.1[500] to 37.48.94.1[500] (1292 bytes)
received packet: from 37.48.94.1[500] to 192.168.1.1[500] (38 bytes)
parsed IKE_SA_INIT response 0 [ N(INVAL_KEY) ]
peer didn't accept DH group MODP_2048, it requested CURVE_25519
initiating IKE_SA PP[1] to 37.48.94.1
generating IKE_SA_INIT request 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) N(FRAG_SUP) N(HASH_ALG) N(RE
DIR_SUP) ]
sending packet: from 192.168.1.1[500] to 37.48.94.1[500] (1068 bytes)
received packet: from 37.48.94.1[500] to 192.168.1.1[500] (265 bytes)
parsed IKE_SA_INIT response 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) CERTREQ N(FRAG_SUP) N(HASH_ALG)
N(CHDLESS_SUP) N(MULT_AUTH) ]
local host is behind NAT, sending keep alives
received cert request for "C=CH, ST=Zug, L=Zug, O=Perfect Privacy, CN=Perfect Privacy IPSEC CA, E=
admin@perfect-privacy.com"
sending cert request for "C=CH, ST=Zug, L=Zug, O=Perfect Privacy, CN=Perfect Privacy IPSEC CA, E=a
```

```
dmin@perfect-privacy.com"
establishing CHILD_SA PP{1}
generating IKE_AUTH request 1 [ IDi CERTREQ CPRQ(ADDR ADDR6 DNS DNS6) SA TSi TSr N(MOBIKE_SUP) N(ADD_4_ADDR) N(ADD_6_ADDR) N(ADD_6_ADDR) N(ADD_6_ADDR) N(MULT_AUTH) N(EAP_ONLY) N(MSG_ID_SYN_SUP) ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (502 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (1248 bytes)
parsed IKE_AUTH response 1 [ EF(1/2) ]
received fragment #1 of 2, waiting for complete IKE message
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (518 bytes)
parsed IKE_AUTH response 1 [ EF(2/2) ]
received fragment #2 of 2, reassembling fragmented IKE message
parsed IKE_AUTH response 1 [ IDr CERT AUTH EAP/REQ/ID ]
received end entity cert "C=CH, O=Perfect Privacy, CN=amsterdam4.perfect-privacy.com"
  using certificate "C=CH, O=Perfect Privacy, CN=amsterdam4.perfect-privacy.com"
  using trusted ca certificate "C=CH, ST=Zug, L=Zug, O=Perfect Privacy, CN=Perfect Privacy IPSEC C
A, E=admin@perfect-privacy.com"
checking certificate status of "C=CH, O=Perfect Privacy, CN=amsterdam4.perfect-privacy.com"
certificate status is not available
  reached self-signed root ca with a path length of 0
authentication of 'amsterdam.perfect-privacy.com' with RSA_EMSA_PKCS1_SHA2_256 successful
server requested EAP_IDENTITY (id 0x00), sending 'Username'
generating IKE_AUTH request 2 [ EAP/RES/ID ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (75 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (97 bytes)
parsed IKE_AUTH response 2 [ EAP/REQ/MSCHAPV2 ]
server requested EAP_MSCHAPV2 authentication (id 0xAC)
generating IKE_AUTH request 3 [ EAP/RES/MSCHAPV2 ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (129 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (134 bytes)
parsed IKE_AUTH response 3 [ EAP/REQ/MSCHAPV2 ]
EAP-MS-CHAPv2 succeeded: 'Welcome2strongSwan'
generating IKE_AUTH request 4 [ EAP/RES/MSCHAPV2 ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (67 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (65 bytes)
parsed IKE_AUTH response 4 [ EAP/SUCC ]
EAP method EAP_MSCHAPV2 succeeded, MSK established
authentication of '192.168.1.1' (myself) with EAP
generating IKE_AUTH request 5 [ AUTH ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (129 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (458 bytes)
parsed IKE_AUTH response 5 [ AUTH CPRP(ADDR ADDR6 DNS DNS) SA TSi TSr N(MOBIKE_SUP) N(ADD_4_ADDR) N(ADD_4_ADDR) N(ADD_4_ADDR) N(ADD_4_ADDR) N(ADD_4_ADDR) N(ADD_4_ADDR) N(ADD_4_ADDR) ]
authentication of 'amsterdam.perfect-privacy.com' with EAP successful
IKE_SA PP[1] established between 192.168.1.1[192.168.1.1]...37.48.94.1[amsterdam.perfect-privacy.com]
installing DNS server 5.79.98.56 to /etc/resolv.conf
installing DNS server 185.17.184.3 to /etc/resolv.conf
installing new virtual IP 10.4.75.139
installing new virtual IP fdbf:1d37:bbe0::68:103:0:39a
CHILD_SA PP{1} established with SPIs cbd24330_i c298e8fd_o and TS 10.4.75.139/32 fdbf:1d37:bbe0::68:103:0:39a/128 === 0.0.0.0/0 ::/0
connection 'PP' established successfully
```

```
root@OpenWrt:~# ipsec statusall
```

```
Status of IKE charon daemon (strongSwan 5.6.3, Linux 4.14.131, armv7l):
  uptime: 5 minutes, since Dec 12 01:25:37 2019
  worker threads: 10 of 16 idle, 6/0/0/0 working, job queue: 0/0/0/0, scheduled: 2
  loaded plugins: charon test-vectors ldap pkcs11 aes des blowfish rc2 sha2 sha1 md4 md5 random no
nce x509 revocation constraints pubkey pkcs1 pkcs7 pkcs8 pkcs12 pgp dnskey sshkey pem openssl gcry
pt af-alg fips-prf gmp curve25519 agent xcbc cmac hmac ctr ccm gcm curl mysql sqlite attr kernel-n
etlink resolve socket-default connmark forecast farp stroke vici smp updown eap-identity eap-md5 e
ap-mschapv2 eap-radius eap-tls xauth-generic xauth-eap dhcp whitelist led duplicheck addrblock uni
ty
Listening IP addresses:
```

```

192.168.1.1
fd31:f82e:665b::1
2a02:908:3030:bc40::1
109.91.76.30
2a02:908:3000:3:18cb:fcf4:f122:9888
Connections:
lan-passthrough: %any...%any IKEv1/2
lan-passthrough: local: uses public key authentication
lan-passthrough: remote: uses public key authentication
lan-passthrough: child: 192.168.1.0/24 === 192.168.1.0/24 PASS
PP: 192.168.1.1...37.48.94.1 IKEv2, dpddelay=300s
PP: local: [192.168.1.1] uses EAP_MSCHAPV2 authentication with EAP identity 'Username
'
PP: remote: uses public key authentication
PP: child: dynamic === 0.0.0.0/0 ::/0 TUNNEL, dpdaction=restart
Shunted Connections:
lan-passthrough: 192.168.1.0/24 === 192.168.1.0/24 PASS
Security Associations (1 up, 0 connecting):
PP[1]: ESTABLISHED 2 minutes ago, 109.91.76.30[192.168.1.1]...37.48.94.1[amsterdam.perfe
ct-privacy.com]
PP[1]: IKEv2 SPIs: 1d66316ba2216c5a_i* bb8f29c0db8008df_r, rekeying disabled
PP[1]: IKE proposal: AES_GCM_16_256/PRF_HMAC_SHA2_512/CURVE_25519
PP{1}: INSTALLED, TUNNEL, reqid 1, ESP in UDP SPIs: cbd24330_i c298e8fd_o
PP{1}: AES_CBC_256/HMAC_SHA2_256_128, 62897 bytes_i (165 pkts, 1s ago), 63968 bytes_o (
228 pkts, 1s ago), rekeying disabled
PP{1}: 10.4.75.139/32 fdbf:1d37:bbe0::68:103:0:39a/128 === 0.0.0.0/0 ::/0

```

```
root@OpenWrt:~# iptables-save
```

```

# Generated by iptables-save v1.6.2 on Thu Dec 12 01:38:26 2019
*nat
:PREROUTING ACCEPT [57:3574]
:INPUT ACCEPT [37:2548]
:OUTPUT ACCEPT [66:4964]
:POSTROUTING ACCEPT [0:0]
:postrouting_lan_rule - [0:0]
:postrouting_rule - [0:0]
:postrouting_wan_rule - [0:0]
:prerouting_lan_rule - [0:0]
:prerouting_rule - [0:0]
:prerouting_wan_rule - [0:0]
:zone_lan_postrouting - [0:0]
:zone_lan_prerouting - [0:0]
:zone_wan_postrouting - [0:0]
:zone_wan_prerouting - [0:0]
-A PREROUTING -m comment --comment "!fw3: Custom prerouting rule chain" -j prerouting_rule
-A PREROUTING -i br-lan -m comment --comment "!fw3" -j zone_lan_prerouting
-A PREROUTING -i eth0.2 -m comment --comment "!fw3" -j zone_wan_prerouting
-A POSTROUTING -m comment --comment "!fw3: Custom postrouting rule chain" -j postrouting_rule
-A POSTROUTING -o br-lan -m comment --comment "!fw3" -j zone_lan_postrouting
-A POSTROUTING -o eth0.2 -m comment --comment "!fw3" -j zone_wan_postrouting
-A postrouting_wan_rule -s 192.168.1.0/24 -m policy --dir out --pol none -j SNAT --to-source 10.4.
75.139
-A zone_lan_postrouting -m comment --comment "!fw3: Custom lan postrouting rule chain" -j postrout
ing_lan_rule
-A zone_lan_prerouting -m comment --comment "!fw3: Custom lan prerouting rule chain" -j prerouting
_lan_rule
-A zone_wan_postrouting -m comment --comment "!fw3: Custom wan postrouting rule chain" -j postrout
ing_wan_rule
-A zone_wan_postrouting -m comment --comment "!fw3" -j MASQUERADE
-A zone_wan_prerouting -m comment --comment "!fw3: Custom wan prerouting rule chain" -j prerouting
_wan_rule
COMMIT
# Completed on Thu Dec 12 01:38:26 2019
# Generated by iptables-save v1.6.2 on Thu Dec 12 01:38:26 2019
*mangle

```

```

:PREROUTING ACCEPT [2097:616262]
:INPUT ACCEPT [1367:295447]
:FORWARD ACCEPT [730:320815]
:OUTPUT ACCEPT [1312:200925]
:POSTROUTING ACCEPT [2039:521620]
-A FORWARD -o eth0.2 -p tcp -m tcp --tcp-flags SYN,RST SYN -m comment --comment "!fw3: Zone wan MT
U fixing" -j TCPMSS --clamp-mss-to-pmtu
COMMIT
# Completed on Thu Dec 12 01:38:26 2019
# Generated by iptables-save v1.6.2 on Thu Dec 12 01:38:26 2019
*filter
:INPUT ACCEPT [0:0]
:FORWARD DROP [0:0]
:OUTPUT ACCEPT [0:0]
:forwarding_lan_rule - [0:0]
:forwarding_rule - [0:0]
:forwarding_wan_rule - [0:0]
:input_lan_rule - [0:0]
:input_rule - [0:0]
:input_wan_rule - [0:0]
:output_lan_rule - [0:0]
:output_rule - [0:0]
:output_wan_rule - [0:0]
:reject - [0:0]
:syn_flood - [0:0]
:zone_lan_dest_ACCEPT - [0:0]
:zone_lan_forward - [0:0]
:zone_lan_input - [0:0]
:zone_lan_output - [0:0]
:zone_lan_src_ACCEPT - [0:0]
:zone_wan_dest_ACCEPT - [0:0]
:zone_wan_dest_REJECT - [0:0]
:zone_wan_forward - [0:0]
:zone_wan_input - [0:0]
:zone_wan_output - [0:0]
:zone_wan_src_REJECT - [0:0]
-A INPUT -d 10.4.75.139/32 -i br-lan -m policy --dir in --pol ipsec --reqid 1 --proto esp -j ACCEP
T
-A INPUT -i lo -m comment --comment "!fw3" -j ACCEPT
-A INPUT -m comment --comment "!fw3: Custom input rule chain" -j input_rule
-A INPUT -m conntrack --ctstate RELATED,ESTABLISHED -m comment --comment "!fw3" -j ACCEPT
-A INPUT -p tcp -m tcp --tcp-flags FIN,SYN,RST,ACK SYN -m comment --comment "!fw3" -j syn_flood
-A INPUT -i br-lan -m comment --comment "!fw3" -j zone_lan_input
-A INPUT -i eth0.2 -m comment --comment "!fw3" -j zone_wan_input
-A FORWARD -d 10.4.75.139/32 -i br-lan -m policy --dir in --pol ipsec --reqid 1 --proto esp -j ACC
EPT
-A FORWARD -s 10.4.75.139/32 -o br-lan -m policy --dir out --pol ipsec --reqid 1 --proto esp -j AC
CEPT
-A FORWARD -m comment --comment "!fw3: Custom forwarding rule chain" -j forwarding_rule
-A FORWARD -m conntrack --ctstate RELATED,ESTABLISHED -m comment --comment "!fw3" -j ACCEPT
-A FORWARD -i br-lan -m comment --comment "!fw3" -j zone_lan_forward
-A FORWARD -i eth0.2 -m comment --comment "!fw3" -j zone_wan_forward
-A FORWARD -m comment --comment "!fw3" -j reject
-A OUTPUT -s 10.4.75.139/32 -o br-lan -m policy --dir out --pol ipsec --reqid 1 --proto esp -j ACC
EPT
-A OUTPUT -o lo -m comment --comment "!fw3" -j ACCEPT
-A OUTPUT -m comment --comment "!fw3: Custom output rule chain" -j output_rule
-A OUTPUT -m conntrack --ctstate RELATED,ESTABLISHED -m comment --comment "!fw3" -j ACCEPT
-A OUTPUT -o br-lan -m comment --comment "!fw3" -j zone_lan_output
-A OUTPUT -o eth0.2 -m comment --comment "!fw3" -j zone_wan_output
-A reject -p tcp -m comment --comment "!fw3" -j REJECT --reject-with tcp-reset
-A reject -m comment --comment "!fw3" -j REJECT --reject-with icmp-port-unreachable
-A syn_flood -p tcp -m tcp --tcp-flags FIN,SYN,RST,ACK SYN -m limit --limit 25/sec --limit-burst 5
0 -m comment --comment "!fw3" -j RETURN
-A syn_flood -m comment --comment "!fw3" -j DROP
-A zone_lan_dest_ACCEPT -o br-lan -m comment --comment "!fw3" -j ACCEPT
-A zone_lan_forward -m comment --comment "!fw3: Custom lan forwarding rule chain" -j forwarding_la

```

```

n_rule
-A zone_lan_forward -m comment --comment "!fw3: Zone lan to wan forwarding policy" -j zone_wan_dest_ACCEPT
-A zone_lan_forward -m conntrack --ctstate DNAT -m comment --comment "!fw3: Accept port forwards" -j ACCEPT
-A zone_lan_forward -m comment --comment "!fw3" -j zone_lan_dest_ACCEPT
-A zone_lan_input -m comment --comment "!fw3: Custom lan input rule chain" -j input_lan_rule
-A zone_lan_input -m conntrack --ctstate DNAT -m comment --comment "!fw3: Accept port redirections" -j ACCEPT
-A zone_lan_input -m comment --comment "!fw3" -j zone_lan_src_ACCEPT
-A zone_lan_output -m comment --comment "!fw3: Custom lan output rule chain" -j output_lan_rule
-A zone_lan_output -m comment --comment "!fw3" -j zone_lan_dest_ACCEPT
-A zone_lan_src_ACCEPT -i br-lan -m conntrack --ctstate NEW,UNTRACKED -m comment --comment "!fw3" -j ACCEPT
-A zone_wan_dest_ACCEPT -o eth0.2 -m conntrack --ctstate INVALID -m comment --comment "!fw3: Prevent NAT leakage" -j DROP
-A zone_wan_dest_ACCEPT -o eth0.2 -m comment --comment "!fw3" -j ACCEPT
-A zone_wan_dest_REJECT -o eth0.2 -m comment --comment "!fw3" -j reject
-A zone_wan_forward -m comment --comment "!fw3: Custom wan forwarding rule chain" -j forwarding_wan_rule
-A zone_wan_forward -p esp -m comment --comment "!fw3: Allow-IPSec-ESP" -j zone_lan_dest_ACCEPT
-A zone_wan_forward -p udp -m udp --dport 500 -m comment --comment "!fw3: Allow-ISAKMP" -j zone_lan_dest_ACCEPT
-A zone_wan_forward -m conntrack --ctstate DNAT -m comment --comment "!fw3: Accept port forwards" -j ACCEPT
-A zone_wan_forward -m comment --comment "!fw3" -j zone_wan_dest_REJECT
-A zone_wan_input -m comment --comment "!fw3: Custom wan input rule chain" -j input_wan_rule
-A zone_wan_input -p udp -m udp --dport 68 -m comment --comment "!fw3: Allow-DHCP-Renew" -j ACCEPT
-A zone_wan_input -p icmp -m icmp --icmp-type 8 -m comment --comment "!fw3: Allow-Ping" -j ACCEPT
-A zone_wan_input -p igmp -m comment --comment "!fw3: Allow-IGMP" -j ACCEPT
-A zone_wan_input -m conntrack --ctstate DNAT -m comment --comment "!fw3: Accept port redirections" -j ACCEPT
-A zone_wan_input -m comment --comment "!fw3" -j zone_wan_src_REJECT
-A zone_wan_output -m comment --comment "!fw3: Custom wan output rule chain" -j output_wan_rule
-A zone_wan_output -m comment --comment "!fw3" -j zone_wan_dest_ACCEPT
-A zone_wan_src_REJECT -i eth0.2 -m comment --comment "!fw3" -j reject
COMMIT
# Completed on Thu Dec 12 01:38:26 2019

```

root@OpenWrt:~# ip6tables-save

```

# Generated by ip6tables-save v1.6.2 on Thu Dec 12 01:39:58 2019
*mangle
:PREROUTING ACCEPT [1111:85007]
:INPUT ACCEPT [967:71246]
:FORWARD ACCEPT [65:6315]
:OUTPUT ACCEPT [182:19984]
:POSTROUTING ACCEPT [247:26299]
-A FORWARD -o eth0.2 -p tcp -m tcp --tcp-flags SYN,RST SYN -m comment --comment "!fw3: Zone wan MTU fixing" -j TCPMSS --clamp-mss-to-pmtu
COMMIT
# Completed on Thu Dec 12 01:39:58 2019
# Generated by ip6tables-save v1.6.2 on Thu Dec 12 01:39:58 2019
*filter
:INPUT ACCEPT [0:0]
:FORWARD DROP [0:0]
:OUTPUT ACCEPT [0:0]
:forwarding_lan_rule - [0:0]
:forwarding_rule - [0:0]
:forwarding_wan_rule - [0:0]
:input_lan_rule - [0:0]
:input_rule - [0:0]
:input_wan_rule - [0:0]
:output_lan_rule - [0:0]
:output_rule - [0:0]
:output_wan_rule - [0:0]

```

```

:reject - [0:0]
:syn_flood - [0:0]
:zone_lan_dest_ACCEPT - [0:0]
:zone_lan_forward - [0:0]
:zone_lan_input - [0:0]
:zone_lan_output - [0:0]
:zone_lan_src_ACCEPT - [0:0]
:zone_wan_dest_ACCEPT - [0:0]
:zone_wan_dest_REJECT - [0:0]
:zone_wan_forward - [0:0]
:zone_wan_input - [0:0]
:zone_wan_output - [0:0]
:zone_wan_src_REJECT - [0:0]
-A INPUT -d fdbf:1d37:bbe0::68:103:0:3e7/128 -i br-lan -m policy --dir in --pol ipsec --reqid 1 --
proto esp -j ACCEPT
-A INPUT -i lo -m comment --comment "!fw3" -j ACCEPT
-A INPUT -m comment --comment "!fw3: Custom input rule chain" -j input_rule
-A INPUT -m conntrack --ctstate RELATED,ESTABLISHED -m comment --comment "!fw3" -j ACCEPT
-A INPUT -p tcp -m tcp --tcp-flags FIN,SYN,RST,ACK SYN -m comment --comment "!fw3" -j syn_flood
-A INPUT -i br-lan -m comment --comment "!fw3" -j zone_lan_input
-A INPUT -i eth0.2 -m comment --comment "!fw3" -j zone_wan_input
-A FORWARD -d fdbf:1d37:bbe0::68:103:0:3e7/128 -i br-lan -m policy --dir in --pol ipsec --reqid 1
--proto esp -j ACCEPT
-A FORWARD -s fdbf:1d37:bbe0::68:103:0:3e7/128 -o br-lan -m policy --dir out --pol ipsec --reqid 1
--proto esp -j ACCEPT
-A FORWARD -m comment --comment "!fw3: Custom forwarding rule chain" -j forwarding_rule
-A FORWARD -m conntrack --ctstate RELATED,ESTABLISHED -m comment --comment "!fw3" -j ACCEPT
-A FORWARD -i br-lan -m comment --comment "!fw3" -j zone_lan_forward
-A FORWARD -i eth0.2 -m comment --comment "!fw3" -j zone_wan_forward
-A FORWARD -m comment --comment "!fw3" -j reject
-A OUTPUT -s fdbf:1d37:bbe0::68:103:0:3e7/128 -o br-lan -m policy --dir out --pol ipsec --reqid 1
--proto esp -j ACCEPT
-A OUTPUT -o lo -m comment --comment "!fw3" -j ACCEPT
-A OUTPUT -m comment --comment "!fw3: Custom output rule chain" -j output_rule
-A OUTPUT -m conntrack --ctstate RELATED,ESTABLISHED -m comment --comment "!fw3" -j ACCEPT
-A OUTPUT -o br-lan -m comment --comment "!fw3" -j zone_lan_output
-A OUTPUT -o eth0.2 -m comment --comment "!fw3" -j zone_wan_output
-A reject -p tcp -m comment --comment "!fw3" -j REJECT --reject-with tcp-reset
-A reject -m comment --comment "!fw3" -j REJECT --reject-with icmp6-port-unreachable
-A syn_flood -p tcp -m tcp --tcp-flags FIN,SYN,RST,ACK SYN -m limit --limit 25/sec --limit-burst 5
0 -m comment --comment "!fw3" -j RETURN
-A syn_flood -m comment --comment "!fw3" -j DROP
-A zone_lan_dest_ACCEPT -o br-lan -m comment --comment "!fw3" -j ACCEPT
-A zone_lan_forward -m comment --comment "!fw3: Custom lan forwarding rule chain" -j forwarding_la
n_rule
-A zone_lan_forward -m comment --comment "!fw3: Zone lan to wan forwarding policy" -j zone_wan_des
t_ACCEPT
-A zone_lan_forward -m comment --comment "!fw3" -j zone_lan_dest_ACCEPT
-A zone_lan_input -m comment --comment "!fw3: Custom lan input rule chain" -j input_lan_rule
-A zone_lan_input -m comment --comment "!fw3" -j zone_lan_src_ACCEPT
-A zone_lan_output -m comment --comment "!fw3: Custom lan output rule chain" -j output_lan_rule
-A zone_lan_output -m comment --comment "!fw3" -j zone_lan_dest_ACCEPT
-A zone_lan_src_ACCEPT -i br-lan -m conntrack --ctstate NEW,UNTRACKED -m comment --comment "!fw3"
-j ACCEPT
-A zone_wan_dest_ACCEPT -o eth0.2 -m conntrack --ctstate INVALID -m comment --comment "!fw3: Preve
nt NAT leakage" -j DROP
-A zone_wan_dest_ACCEPT -o eth0.2 -m comment --comment "!fw3" -j ACCEPT
-A zone_wan_dest_REJECT -o eth0.2 -m comment --comment "!fw3" -j reject
-A zone_wan_forward -m comment --comment "!fw3: Custom wan forwarding rule chain" -j forwarding_wa
n_rule
-A zone_wan_forward -p ipv6-icmp -m icmp6 --icmpv6-type 128 -m limit --limit 1000/sec -m comment -
comment "!fw3: Allow-ICMPv6-Forward" -j ACCEPT
-A zone_wan_forward -p ipv6-icmp -m icmp6 --icmpv6-type 129 -m limit --limit 1000/sec -m comment -
comment "!fw3: Allow-ICMPv6-Forward" -j ACCEPT
-A zone_wan_forward -p ipv6-icmp -m icmp6 --icmpv6-type 1 -m limit --limit 1000/sec -m comment --c
omment "!fw3: Allow-ICMPv6-Forward" -j ACCEPT
-A zone_wan_forward -p ipv6-icmp -m icmp6 --icmpv6-type 2 -m limit --limit 1000/sec -m comment --c

```

```

omment "!fw3: Allow-ICMPv6-Forward" -j ACCEPT
-A zone_wan_forward -p ipv6-icmp -m icmp6 --icmpv6-type 3 -m limit --limit 1000/sec -m comment --c
omment "!fw3: Allow-ICMPv6-Forward" -j ACCEPT
-A zone_wan_forward -p ipv6-icmp -m icmp6 --icmpv6-type 4/0 -m limit --limit 1000/sec -m comment -
-comment "!fw3: Allow-ICMPv6-Forward" -j ACCEPT
-A zone_wan_forward -p ipv6-icmp -m icmp6 --icmpv6-type 4/1 -m limit --limit 1000/sec -m comment -
-comment "!fw3: Allow-ICMPv6-Forward" -j ACCEPT
-A zone_wan_forward -p esp -m comment --comment "!fw3: Allow-IPSec-ESP" -j zone_lan_dest_ACCEPT
-A zone_wan_forward -p udp -m udp --dport 500 -m comment --comment "!fw3: Allow-ISAKMP" -j zone_la
n_dest_ACCEPT
-A zone_wan_forward -m comment --comment "!fw3" -j zone_wan_dest_REJECT
-A zone_wan_input -m comment --comment "!fw3: Custom wan input rule chain" -j input_wan_rule
-A zone_wan_input -s fc00::/6 -d fc00::/6 -p udp -m udp --dport 546 -m comment --comment "!fw3: Al
low-DHCPv6" -j ACCEPT
-A zone_wan_input -s fe80::/10 -p ipv6-icmp -m icmp6 --icmpv6-type 130/0 -m comment --comment "!fw
3: Allow-MLD" -j ACCEPT
-A zone_wan_input -s fe80::/10 -p ipv6-icmp -m icmp6 --icmpv6-type 131/0 -m comment --comment "!fw
3: Allow-MLD" -j ACCEPT
-A zone_wan_input -s fe80::/10 -p ipv6-icmp -m icmp6 --icmpv6-type 132/0 -m comment --comment "!fw
3: Allow-MLD" -j ACCEPT
-A zone_wan_input -s fe80::/10 -p ipv6-icmp -m icmp6 --icmpv6-type 143/0 -m comment --comment "!fw
3: Allow-MLD" -j ACCEPT
-A zone_wan_input -p ipv6-icmp -m icmp6 --icmpv6-type 128 -m limit --limit 1000/sec -m comment --c
omment "!fw3: Allow-ICMPv6-Input" -j ACCEPT
-A zone_wan_input -p ipv6-icmp -m icmp6 --icmpv6-type 129 -m limit --limit 1000/sec -m comment --c
omment "!fw3: Allow-ICMPv6-Input" -j ACCEPT
-A zone_wan_input -p ipv6-icmp -m icmp6 --icmpv6-type 1 -m limit --limit 1000/sec -m comment --com
ment "!fw3: Allow-ICMPv6-Input" -j ACCEPT
-A zone_wan_input -p ipv6-icmp -m icmp6 --icmpv6-type 2 -m limit --limit 1000/sec -m comment --com
ment "!fw3: Allow-ICMPv6-Input" -j ACCEPT
-A zone_wan_input -p ipv6-icmp -m icmp6 --icmpv6-type 3 -m limit --limit 1000/sec -m comment --com
ment "!fw3: Allow-ICMPv6-Input" -j ACCEPT
-A zone_wan_input -p ipv6-icmp -m icmp6 --icmpv6-type 4/0 -m limit --limit 1000/sec -m comment --c
omment "!fw3: Allow-ICMPv6-Input" -j ACCEPT
-A zone_wan_input -p ipv6-icmp -m icmp6 --icmpv6-type 4/1 -m limit --limit 1000/sec -m comment --c
omment "!fw3: Allow-ICMPv6-Input" -j ACCEPT
-A zone_wan_input -p ipv6-icmp -m icmp6 --icmpv6-type 133 -m limit --limit 1000/sec -m comment --c
omment "!fw3: Allow-ICMPv6-Input" -j ACCEPT
-A zone_wan_input -p ipv6-icmp -m icmp6 --icmpv6-type 135 -m limit --limit 1000/sec -m comment --c
omment "!fw3: Allow-ICMPv6-Input" -j ACCEPT
-A zone_wan_input -p ipv6-icmp -m icmp6 --icmpv6-type 134 -m limit --limit 1000/sec -m comment --c
omment "!fw3: Allow-ICMPv6-Input" -j ACCEPT
-A zone_wan_input -p ipv6-icmp -m icmp6 --icmpv6-type 136 -m limit --limit 1000/sec -m comment --c
omment "!fw3: Allow-ICMPv6-Input" -j ACCEPT
-A zone_wan_input -m comment --comment "!fw3" -j zone_wan_src_REJECT
-A zone_wan_output -m comment --comment "!fw3: Custom wan output rule chain" -j output_wan_rule
-A zone_wan_output -m comment --comment "!fw3" -j zone_wan_dest_ACCEPT
-A zone_wan_src_REJECT -i eth0.2 -m comment --comment "!fw3" -j reject
COMMIT
# Completed on Thu Dec 12 01:39:58 2019

```

```
root@OpenWrt:~# ip route show table all
```

```

default via 109.91.76.1 dev eth0.2 table 220 proto static src 10.4.75.139
192.168.1.0/24 dev br-lan table 220 proto static src 192.168.1.1
default via 109.91.76.1 dev eth0.2 proto static src 109.91.76.30
109.91.76.0/22 dev eth0.2 proto kernel scope link src 109.91.76.30
192.168.1.0/24 dev br-lan proto kernel scope link src 192.168.1.1
local 10.4.75.139 dev eth0.2 table local proto kernel scope host src 10.4.75.139
broadcast 109.91.76.0 dev eth0.2 table local proto kernel scope link src 109.91.76.30
local 109.91.76.30 dev eth0.2 table local proto kernel scope host src 109.91.76.30
broadcast 109.91.79.255 dev eth0.2 table local proto kernel scope link src 109.91.76.30
broadcast 127.0.0.0 dev lo table local proto kernel scope link src 127.0.0.1
local 127.0.0.0/8 dev lo table local proto kernel scope host src 127.0.0.1
local 127.0.0.1 dev lo table local proto kernel scope host src 127.0.0.1
broadcast 127.255.255.255 dev lo table local proto kernel scope link src 127.0.0.1

```



```

broadcast 192.168.1.0 dev br-lan table local proto kernel scope link src 192.168.1.1
local 192.168.1.1 dev br-lan table local proto kernel scope host src 192.168.1.1
broadcast 192.168.1.255 dev br-lan table local proto kernel scope link src 192.168.1.1
default dev eth0.2 table 220 proto static metric 1024 pref medium
default from 2a02:908:3000:3:18cb:fcf4:f122:9888 via fe80::201:5cff:fe92:9846 dev eth0.2 proto sta
tic metric 512 pref medium
default from 2a02:908:3030:bc40::/59 via fe80::201:5cff:fe92:9846 dev eth0.2 proto static metric 5
12 pref medium
2a02:908:3030:bc40::/64 dev br-lan proto static metric 1024 pref medium
unreachable 2a02:908:3030:bc40::/59 dev lo proto static metric 2147483647 error -113 pref medium
fd31:f82e:665b::/64 dev br-lan proto static metric 1024 pref medium
unreachable fd31:f82e:665b::/48 dev lo proto static metric 2147483647 error -113 pref medium
fdbf:1d37:bbe0::68:103:0:3e7 dev eth0.2 proto kernel metric 256 pref medium
fe80::/64 dev eth0 proto kernel metric 256 pref medium
fe80::/64 dev eth0.2 proto kernel metric 256 pref medium
fe80::/64 dev br-lan proto kernel metric 256 pref medium
local ::1 dev lo table local proto kernel metric 0 pref medium
local 2a02:908:3000:3:18cb:fcf4:f122:9888 dev eth0.2 table local proto kernel metric 0 pref medium
anycast 2a02:908:3030:bc40:: dev br-lan table local proto kernel metric 0 pref medium
local 2a02:908:3030:bc40::1 dev br-lan table local proto kernel metric 0 pref medium
anycast fd31:f82e:665b:: dev br-lan table local proto kernel metric 0 pref medium
local fd31:f82e:665b::1 dev br-lan table local proto kernel metric 0 pref medium
local fdbf:1d37:bbe0::68:103:0:3e7 dev eth0.2 table local proto kernel metric 0 pref medium
anycast fe80:: dev eth0.2 table local proto kernel metric 0 pref medium
anycast fe80:: dev eth0 table local proto kernel metric 0 pref medium
anycast fe80:: dev br-lan table local proto kernel metric 0 pref medium
local fe80::da50:e6ff:fe4f:9848 dev eth0 table local proto kernel metric 0 pref medium
local fe80::da50:e6ff:fe4f:9848 dev br-lan table local proto kernel metric 0 pref medium
local fe80::da50:e6ff:fe4f:9849 dev eth0.2 table local proto kernel metric 0 pref medium
ff00::/8 dev eth0 table local metric 256 pref medium
ff00::/8 dev br-lan table local metric 256 pref medium
ff00::/8 dev eth0.2 table local metric 256 pref medium

```

root@OpenWrt:~# ip address

```

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,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 100
0
    link/ether d8:50:e6:4f:98:48 brd ff:ff:ff:ff:ff:ff
    inet6 fe80::da50:e6ff:fe4f:9848/64 scope link
        valid_lft forever preferred_lft forever
3: eth1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
    link/ether 6e:a3:74:5b:45:50 brd ff:ff:ff:ff:ff:ff
4: eth2: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
    link/ether 3e:c6:14:ce:91:28 brd ff:ff:ff:ff:ff:ff
6: ip_vti0@NONE: <NOARP> mtu 1480 qdisc noop state DOWN group default qlen 1000
    link/ipip 0.0.0.0 brd 0.0.0.0
7: wlan0: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
    link/ether d8:50:e6:4f:98:48 brd ff:ff:ff:ff:ff:ff
8: br-lan: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 10
00
    link/ether d8:50:e6:4f:98:48 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.1/24 brd 192.168.1.255 scope global br-lan
        valid_lft forever preferred_lft forever
    inet6 2a02:908:3030:bc40::1/60 scope global dynamic noprefixroute
        valid_lft 1209060sec preferred_lft 604260sec
    inet6 fd31:f82e:665b::1/60 scope global noprefixroute
        valid_lft forever preferred_lft forever
    inet6 fe80::da50:e6ff:fe4f:9848/64 scope link
        valid_lft forever preferred_lft forever
9: eth0.1@eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue master br-lan state UP gr

```

```

oup default qlen 1000
  link/ether d8:50:e6:4f:98:48 brd ff:ff:ff:ff:ff:ff
10: eth0.2@eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default q
len 1000
  link/ether d8:50:e6:4f:98:49 brd ff:ff:ff:ff:ff:ff
  inet 109.91.76.30/22 brd 109.91.79.255 scope global eth0.2
    valid_lft forever preferred_lft forever
  inet 10.4.75.139/32 scope global eth0.2
    valid_lft forever preferred_lft forever
  inet6 fdbf:1d37:bbe0::68:103:0:3e7/128 scope global nodad deprecated
    valid_lft forever preferred_lft 0sec
  inet6 2a02:908:3000:3:18cb:fcf4:f122:9888/128 scope global dynamic noprefixroute
    valid_lft 1209061sec preferred_lft 604261sec
  inet6 fe80::da50:e6ff:fe4f:9849/64 scope link
    valid_lft forever preferred_lft forever

```

After the IKEv2 connection:

The IKEv2 VPN IPv4 IP is displayed.

However, the IKEv2 VPN IPv6 IP is not displayed at Check-IP, but my VPN provider offers an IKEv2 IPv6. I get an answer at Check-IP:

```
You don't seem to have an IPv6 capable connection.
```

1) Have I configured ipsec.conf correctly? Are there any suggestions for improving the configuration?

I don't have enough experience, but I don't think OpenWRT can handle IKEv2 IPv6. If it does work, I'd like the solution.

I disabled IPv6 afterwards with:

/etc/sysctl.conf:

```

net.ipv6.conf.all.disable_ipv6 = 1
net.ipv6.conf.default.disable_ipv6 = 1
net.ipv6.conf.lo.disable_ipv6 = 1

```

2) Problem: The content of sysctl.conf completely disables IPv6.

Is there an alternative? I only want to disable IPv6 during the IKEv2 connection.

If everything is configured correctly, then I proceeded as follows:

- I have added an ipsec0 interface, to add a KillSwitch via WebUI and firewall.

- With this configuration:

add to /etc/ipsec.conf:

```

mark_in=42
mark_out=42

```

/etc/strongswan.conf:

```

# strongswan.conf - strongSwan configuration file
#
# Refer to the strongswan.conf(5) manpage for details
#
# Configuration changes should be made in the included files

```

```

charon {
  install_routes=no
  install_virtual_ip=no

  load_modular = yes
  plugins {
    include strongswan.d/charon/*.conf
  }
}

```

```
include strongswan.d/*.conf
```

Terminal:

```
ip tunnel add ipsec0 local 192.168.1.1 remote 37.48.94.1 mode vti key 42
sysctl -w net.ipv4.conf.ipsec0.disable_policy=1
ip link set ipsec0 up
ip route add 10.0.0.0/24 dev ipsec0
ifconfig ipsec0 192.168.1.1 netmask 255.255.255.0 broadcast 192.168.1.255
```

3) Did I add an ipsec0 interface correctly?

All strongSwan IKEv2 settings run without the firewall rules of OpenWRT.

I can not manage the ipsec0 interface through the firewall rules of OpenWRT WebUI.

If there's a KillSwitch for strongSwan, I'd like to know, how to configure it without OpenWRT.

Tell me if you need any more logs.

Best regards

Bernd

History

#1 - 12.12.2019 11:30 - Tobias Brunner

- Status changed from New to Feedback

Your updown script obviously would have to be adapted to handle the virtual IPv6 address and IPv6 traffic (if the NAT is essential you'll have to do the same for IPv6).

If you can't make it work, don't negotiate an IPv6 traffic selector (*rightsubnet*) and don't request a virtual IPv6 address (*leftsourceip*) and then block all IPv6 traffic via your firewall or IPsec block policies (could also be done via updown script if you only want to do that while a connection is established, not really the purpose of a kill switch, though).

If everything is configured correctly, then I proceeded as follows:

- I have added an ipsec0 interface, to add a KillSwitch via WebUI and firewall.

Doesn't seem necessary, but why not complicate things.

I can not manage the ipsec0 interface through the firewall rules of OpenWRT WebUI.

Not really surprising, that's a virtual VTI device.

If there's a KillSwitch for strongSwan, I'd like to know, how to configure it without OpenWRT.

You can do that with IPsec drop policies (similar to [passthrough policies](#), just a different type), or via firewall rules (e.g. with a default DROP policy, the default updown script can then install specific rules to allow tunneled traffic).

#2 - 13.12.2019 23:45 - Bernd Bernikov

Tobias Brunner wrote:

Your updown script obviously would have to be adapted to handle the virtual IPv6 address and IPv6 traffic

At the moment I can't get the updown script from the VPN provider. So I removed from "leftsourceip" %config6 and from "rightsubnet" ::/0.

After that I had an IPv6 leak. That's why I disabled IPv6 with these commandos:

/etc/sysctl.conf:

```
net.ipv6.conf.all.disable_ipv6 = 1
net.ipv6.conf.default.disable_ipv6 = 1
net.ipv6.conf.lo.disable_ipv6 = 1
```

If everything is configured correctly, then I proceeded as follows:

- I have added an ipsec0 interface, to add a KillSwitch via WebUI and firewall.

Doesn't seem necessary, but why not complicate things.

I can not manage the ipsec0 interface through the firewall rules of OpenWRT WebUI.

Not really surprising, that's a virtual VTI device.

Is there an easier way to add an ipsec interface?

If there's a KillSwitch for strongSwan, I'd like to know, how to configure it without OpenWRT.

You can do that with IPsec drop policies (similar to passthrough policies, just a different type)

Thanks for the link. I would like to try this way, but as a beginner I still don't know how to do this.

or via firewall rules

You mean with iptables? Is there an example somewhere?

#3 - 16.12.2019 17:49 - Bernd Bernikov

@Tobias Brunner

Can you help?

#4 - 17.12.2019 15:12 - Bernd Bernikov

Dear strongSwan administrators,

can you recommend a forum that helps me with the issue "IPSec drop policies"?

Greetings

Bernd

#5 - 18.12.2019 10:35 - Tobias Brunner

As mentioned, they are basically like passthrough policies, they just drop the matching packets and they always have a lower priority than other policies (unless the priority is manually changed). See the [support page](#) if you need more help.

#6 - 10.02.2020 03:15 - Bernd Bernikov

Hello.

I may have found a way for kill switch. It seems to work.

But is this method safe or unsafe?

```
connections {
    dropall {
        children {
            dropall {
                local_ts = 0.0.0.0/0
                remote_ts = 0.0.0.0/0
                priority = 2
                mode = drop
                start_action = trap
            }
        }
    }
    lan-passthrough {
        children {
            lan-passthrough {
                local_ts = 192.168.1.0/24 # Replace with your LAN subnet
                remote_ts = 192.168.1.0/24 # Replace with your LAN subnet
                priority = 1
                mode = pass
                start_action = trap
            }
        }
    }
    pp {
        unique = never
    }
}
```

```

version = 2
keyingtries=0
dpd_delay = 300s
rekey_time = 0
encap = yes
proposals = aes256-sha256-modp2048
vips = 0.0.0.0
send_cert = never
send_certreq = yes
local_addrs = 192.168.1.1 # Replace with your default Router IP address
remote_addrs = <PP Server IP> # Replace with your PP Server IP

local {
    id = 192.168.1.1 # Replace with your default Router IP address
    auth = eap-mschapv2
    eap_id = Username # Replace with your PP-Username
}
remote {
    id = %any
    auth = pubkey
}
children {
    pp {
        dpd_action = start
        close_action = start
        inactivity = 36000s
        life_time = 0
        esp_proposals = aes256-sha256
        updown = /etc/swanctl/updown.sh
        remote_ts = 0.0.0.0/0
        priority = 1
        mode = tunnel
        start_action = start # "none" is for manual start, or use "start" for autostart
    }
}
} # connections
secrets {
    eap-user {
        id = Username # Replace with your PP-Username
        secret = "Password" # Replace with your "PP-Password"
    }
} # secrets

```

/etc/swanctl/updown.sh:

```

#!/bin/sh

case "$PLUTO_VERB" in
up-client)
    iptables -t nat -A postrouting_wan_rule -s 192.168.1.0/24 -m policy --dir out --pol none -j SNAT --to-source "$PLUTO_MY_SOURCEIP4_1"
    swanctl --uninstall --child dropall
    ;;
down-client)
    swanctl --install --child dropall
    iptables -t nat -F postrouting_wan_rule
    ;;
esac

```

Best regards.

Bernd

Edit:

It's not safe. There is an ip-leak somewhere between start and updown.sh start for 2 seconds. I don't know where exactly the ip-leak is caused.

I use this configuration for startup (strongswan.conf):

```

# strongswan.conf - strongSwan configuration file
#
# Refer to the strongswan.conf(5) manpage for details
#

```

```
# Configuration changes should be made in the included files
```

```
charon {  
    load_modular = yes  
    plugins {  
        include strongswan.d/charon/*.conf  
    }  
    start-scripts {  
        load-all = /usr/sbin/swanctl --load-all  
    }  
}
```

```
include strongswan.d/*.conf
```

Any idea how to solve this?

#7 - 10.02.2020 11:03 - Tobias Brunner

There is an ip-leak somewhere between start and updown.sh start for 2 seconds.

What does that mean exactly?

Any idea how to solve this?

Why uninstall/install the drop policies in the updown script?

#8 - 10.02.2020 11:28 - Bernd Bernikow

Hello.

What does that mean exactly?

When I restart the router, then I see the real ip from the Internet Service Provider for 2-5 seconds.

Why uninstall/install the drop policies in the updown script?

Because I don't know how else to create the kill switch protection.

Dropall rule block everything. When the IPsec/IKEv2 connection is established, the dropall rule will be deleted (swanctl --uninstall --child dropall). When the IPsec/IKEv2 connection is terminated, dropall rule will be enabled (swanctl --install --child dropall).

Is there an easier kill switch?

Best regards

Bernd

#9 - 10.02.2020 12:44 - Tobias Brunner

When I restart the router, then I see the real ip from the Internet Service Provider for 2-5 seconds.

I guess you'd have to prevent any traffic until the drop policy is installed (will take a while until the daemon is started, the config is loaded and the drop policy is installed in the kernel), e.g. via firewall (or you install the drop policy manually as early as possible). Or do everything via firewall (see below).

Why uninstall/install the drop policies in the updown script?

Because I don't know how else to create the kill switch protection.

You don't have to disable the drop policies if the priority is lower or equal to the actual IPsec policies. Only traffic not matching the IPsec policies will be dropped.

Is there an easier kill switch?

Not sure if it's easier, but you could do this via iptables (default DROP policy and either use an IPsec-policy catch-all rule or install more specific rules

dynamically via updown script).

#10 - 11.02.2020 13:00 - Bernd Bernikov

Hello.

You don't have to disable the drop policies if the priority is lower or equal to the actual IPsec policies. Only traffic not matching the IPsec policies will be dropped.

I know that now. Unfortunately I don't know (still) how the higher rules have to be than dropall rules.

Not sure if it's easier, but you could do this via iptables (default DROP policy and either use an IPsec-policy catch-all rule

iptables DROP all rule I can do, but I don't know what the IPsec-policy catch-all should look like.

Can you help me with catch-all rules? I would have to translate the rules for OpenWrt, because I can't adopt them 1:1.

Best regards

Bernd

#11 - 11.02.2020 18:33 - Tobias Brunner

Unfortunately I don't know (still) how the higher rules have to be than dropall rules.

What do you mean?

Not sure if it's easier, but you could do this via iptables (default DROP policy and either use an IPsec-policy catch-all rule

iptables DROP all rule I can do, but I don't know what the IPsec-policy catch-all should look like.

Can you help me with catch-all rules? I would have to translate the rules for OpenWrt, because I can't adopt them 1:1.

Not sure about OpenWRT, but just use the *policy* module to match packets for which an IPsec policy exists (see our default updown script, I guess you could also just use that one, and the iptables-extensions man page for more).

#12 - 11.02.2020 20:47 - Bernd Bernikov

Hello.

You don't have to disable the drop policies if the priority is lower or equal to the actual IPsec policies. Only traffic not matching the IPsec policies will be dropped.

Okay. Then I didn't understand it. Because in my example the drop policies have the highest priority (2). They are not lower or equal, right?

What do I need to change to not disable drop policies?

Best regards

Bernd

#13 - 12.02.2020 10:39 - Tobias Brunner

Because in my example the drop policies have the highest priority (2). They are not lower or equal, right?

No, it's the lowest (higher number = lower priority). But in your example you don't have to configure any priorities manually, should work fine with the defaults (drop policies have lower priority than passthrough or actual IPsec policies).

What do I need to change to not disable drop policies?

Nothing really, the drop policies have the lowest priorities in your example so they only apply to packets that don't match the passthrough and, if established, the tunnel policies.

#14 - 13.02.2020 02:25 - Bernd Bernikov

Hello.

This are good news for me, if the priorities are ok and if I don't need to disable drop policies.

I am connected to the router via LAN.

If I execute the command "swanctl --load-all" and "swanctl --initiate --child pp" manually, then the IKEv2 connection is established successfully. The connection from LAN to router remains enabled, but I still have no connection from LAN to the Internet.

So is there anything missing?

Best regards

Bernd

#15 - 13.02.2020 09:24 - Tobias Brunner

If I execute the command "swanctl --load-all" and "swanctl --initiate --child pp" manually, then the IKEv2 connection is established successfully. The connection from LAN to router remains enabled, but I still have no connection from LAN to the Internet.

So is there anything missing?

Hard to say. Check traffic counters of IPsec policies/SAs and firewall rules to see where packets go or get dropped. Traffic captures could help too.

#16 - 28.07.2020 03:51 - Bernd Bernikov

Hello.

I've changed my OpenWRT configuration a little. It still receives packets but does not transmit any.

This is the configuration:

/etc/config/network

```
config interface 'ipsec'
    option ifname 'vti0'
    option proto 'static'
    option ipaddr '10.0.0.0'
    option netmask '255.0.0.0'
```

```
config route
    option target '10.0.0.0'
    option netmask '255.0.0.0'
    option interface 'ipsec'
```

/etc/strongswan.conf

```
# strongswan.conf - strongSwan configuration file
#
# Refer to the strongswan.conf(5) manpage for details
#
# Configuration changes should be made in the included files
```

```
charon {
    install_routes = no
    install_virtual_ip = no

    load_modular = yes
    plugins {
        include strongswan.d/charon/*.conf
    }
}
```

/etc/ipsec.conf

```
config setup
    charondebug="all"
    uniqueids=never
```



```

conn PP
    eap_identity="username"
    type=tunnel
    mobike=no
    keyexchange=ikev2
    keyingtries=%forever
    dpdaction=restart
    closeaction=restart
    dpddelay=300s
    inactivity=36000s
    rekey=no
    forceencaps=yes
    authby=secret
    ike=aes256-sha256-modp2048
    esp=aes256-sha256
    leftfirewall=yes
    left=192.168.1.1
    leftid=192.168.1.1
    leftsubnet=10.0.0.0/8
    leftsourceip=%config4
    leftsendcert=never
    leftauth=eap-mschapv2
    rightfirewall=yes
    rightauth=pubkey
    right=37.48.94.1
    rightid=%any
    rightsubnet=0.0.0.0/0
    rightsendcert=always
    leftikeport=4500
    rightikeport=4500
    mark=%unique
    auto=add

```

/etc/ipsec.user

```

# This file is interpreted as shell script.
# Put your custom ip rules here, they will
# be executed with each call to the script
# /usr/lib/ipsec/_updown which by default
# strongswan executes.

# Force it to use vti0
VTI_IF="vti0"

# Private subnet
PRIVATE_SUBNET="192.168.1.0/24"

LOCAL_IF="${PLUTO_INTERFACE}"

# GCP's MTU is 1460, so it's hardcoded
GCP_MTU="1460"

# ipsec overhead is 73 bytes, we need to compute new mtu.
VTI_MTU=$((GCP_MTU-73))

case "${PLUTO_VERB}" in
up-client)
    iptables -t nat -A postrouting_wan_rule -s "${PRIVATE_SUBNET}" -m policy --dir out --pol none -j SNAT --to
-source "${PLUTO_MY_SOURCEIP4_1}"
    ip tunnel add "${VTI_IF}" local "${PLUTO_ME}" remote "${PLUTO_PEER}" mode vti okey "${PLUTO_MARK_OUT%/*}"
    ikey "${PLUTO_MARK_IN%/*}"
    ip link set "${VTI_IF}" up mtu "${VTI_MTU}"
    # Disable policy checks for this interface, otherwise the Kernel will drop the traffic after decryption.
    sysctl -w "net.ipv4.conf.${VTI_IF}.disable_policy=1"
    # Disable RP filter for the tunnel interface
    sysctl -w "net.ipv4.conf.${VTI_IF}.rp_filter=0"
    ;;
down-client)
    iptables -t nat -F postrouting_wan_rule
    ip tunnel del "${VTI_IF}"
    ;;
esac

# Disable IPSEC Encryption on local net

```

```

sysctl -w "net.ipv4.conf.${LOCAL_IF}.disable_policy=1"

root@OpenWrt:~# ipsec up PP
initiating IKE_SA PP[1] to 37.48.94.1
generating IKE_SA_INIT request 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) N(FRAG_SUP) N(HASH_ALG) N(REDIR_SUP) ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (1316 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (38 bytes)
parsed IKE_SA_INIT response 0 [ N(INVAL_KEY) ]
peer didn't accept DH group MODP_2048, it requested CURVE_25519
initiating IKE_SA PP[1] to 37.48.94.1
generating IKE_SA_INIT request 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) N(FRAG_SUP) N(HASH_ALG) N(REDIR_SUP) ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (1092 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (265 bytes)
parsed IKE_SA_INIT response 0 [ SA KE No N(NATD_S_IP) N(NATD_D_IP) CERTREQ N(FRAG_SUP) N(HASH_ALG) N(CHDLESS_SUP) N(MULT_AUTH) ]
selected proposal: IKE:AES_GCM_16_256/PRF_HMAC_SHA2_512/CURVE_25519
local host is behind NAT, sending keep alives
received cert request for "C=CH, ST=Zug, L=Zug, O=Perfect Privacy, CN=Perfect Privacy IPSEC CA, E=admin@perfect-privacy.com"
sending cert request for "C=CH, ST=Zug, L=Zug, O=Perfect Privacy, CN=Perfect Privacy IPSEC CA, E=admin@perfect-privacy.com"
establishing CHILD_SA PP{1}
generating IKE_AUTH request 1 [ IDi CERTREQ CPRQ(ADDR DNS) SA TSi TSr N(MULT_AUTH) N(EAP_ONLY) N(MSG_ID_SYN_SU P) ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (322 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (1248 bytes)
parsed IKE_AUTH response 1 [ EF(1/2) ]
received fragment #1 of 2, waiting for complete IKE message
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (518 bytes)
parsed IKE_AUTH response 1 [ EF(2/2) ]
received fragment #2 of 2, reassembled fragmented IKE message (1701 bytes)
parsed IKE_AUTH response 1 [ IDr CERT AUTH EAP/REQ/ID ]
received end entity cert "C=CH, O=Perfect Privacy, CN=amsterdam4.perfect-privacy.com"
  using certificate "C=CH, O=Perfect Privacy, CN=amsterdam4.perfect-privacy.com"
  using trusted ca certificate "C=CH, ST=Zug, L=Zug, O=Perfect Privacy, CN=Perfect Privacy IPSEC CA, E=admin@perfect-privacy.com"
checking certificate status of "C=CH, O=Perfect Privacy, CN=amsterdam4.perfect-privacy.com"
certificate status is not available
  reached self-signed root ca with a path length of 0
authentication of 'amsterdam.perfect-privacy.com' with RSA_EMSA_PKCS1_SHA2_256 successful
server requested EAP_IDENTITY (id 0x00), sending 'username'
generating IKE_AUTH request 2 [ EAP/RES/ID ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (75 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (97 bytes)
parsed IKE_AUTH response 2 [ EAP/REQ/MSCHAPV2 ]
server requested EAP_MSCHAPV2 authentication (id 0xAB)
generating IKE_AUTH request 3 [ EAP/RES/MSCHAPV2 ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (129 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (134 bytes)
parsed IKE_AUTH response 3 [ EAP/REQ/MSCHAPV2 ]
EAP-MS-CHAPv2 succeeded: 'Welcome2strongSwan'
generating IKE_AUTH request 4 [ EAP/RES/MSCHAPV2 ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (67 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (65 bytes)
parsed IKE_AUTH response 4 [ EAP/SUCC ]
EAP method EAP_MSCHAPV2 succeeded, MSK established
authentication of '192.168.1.1' (myself) with EAP
generating IKE_AUTH request 5 [ AUTH ]
sending packet: from 192.168.1.1[4500] to 37.48.94.1[4500] (129 bytes)
received packet: from 37.48.94.1[4500] to 192.168.1.1[4500] (253 bytes)
parsed IKE_AUTH response 5 [ AUTH CPRP(ADDR DNS DNS) SA TSi TSr ]
authentication of 'amsterdam.perfect-privacy.com' with EAP successful
IKE_SA PP[1] established between 192.168.1.1[192.168.1.1]...37.48.94.1[amsterdam.perfect-privacy.com]
installing DNS server 37.48.94.55 to /etc/resolv.conf
installing DNS server 185.17.184.3 to /etc/resolv.conf
installing new virtual IP 10.4.75.154
selected proposal: ESP:AES_CBC_256/HMAC_SHA2_256_128/NO_EXT_SEQ
CHILD_SA PP{1} established with SPIs call0860_i c66085d8_o and TS 10.4.75.154/32 == 0.0.0.0/0
updown: RTNETLINK answers: File exists
updown: net.ipv4.conf.vti0.disable_policy = 1
updown: net.ipv4.conf.vti0.rp_filter = 0
updown: net.ipv4.conf.eth0.2.disable_policy = 1
connection 'PP' established successfully

root@OpenWrt:~# ipsec statusall

```

```
Status of IKE charon daemon (strongSwan 5.8.2, Linux 4.14.180, armv7l):
  uptime: 3 minutes, since Jul 27 04:22:44 2020
  worker threads: 10 of 16 idle, 6/0/0/0 working, job queue: 0/0/0/0, scheduled: 3
  loaded plugins: charon test-vectors ldap pkcs11 aes des blowfish rc2 sha2 sha1 md4 md5 random nonce x509 revocation constraints pubkey pkcs1 pkcs7 pkcs8 pkcs12 pgp dnskey sshkey pem openssl gcrypt af-alg fips-prf gmp curve25519 agent xcbc cmac hmac ctr ccm gcm curl mysql sqlite attr kernel-netlink resolve socket-default connmark forecast farp stroke vici smp updown eap-identity eap-md5 eap-mschapv2 eap-radius eap-tls xauth-generic xauth-eap dhcp whitelist led duplicheck addrblock unity
Listening IP addresses:
  192.168.1.1
  2a02:908:3033:d220::1
  fd32:998d:b371::1
  62.143.72.48
  2a02:908:3000:3:3081:5087:204:43e7
  10.0.0.0
Connections:
  PP: 192.168.1.1...37.48.94.1 IKEv2, dpddelay=300s
  PP: local: [192.168.1.1] uses EAP_MSCHAPV2 authentication with EAP identity 'username'
  PP: remote: uses public key authentication
  PP: child: 10.0.0.0/8 === 0.0.0.0/0 TUNNEL, dpdaction=restart
Security Associations (1 up, 0 connecting):
  PP[1]: ESTABLISHED 3 minutes ago, 192.168.1.1[192.168.1.1]...37.48.94.1[amsterdam.perfect-privacy.com]
  PP[1]: IKEv2 SPIs: dca18e8dee561326_i* 85c335d8e476a2eb_r, rekeying disabled
  PP[1]: IKE proposal: AES_GCM_16_256/PRF_HMAC_SHA2_512/CURVE_25519
  PP{1}: INSTALLED, TUNNEL, reqid 1, ESP in UDP SPIs: call0860_i c66085d8_o
  PP{1}: AES_CBC_256/HMAC_SHA2_256_128, 60293 bytes_i (163 pkts, 4s ago), 53516 bytes_o (159 pkts, 4s ago), rekeying disabled
  PP{1}: 10.4.75.154/32 === 0.0.0.0/0
```

```
root@OpenWrt:~# ip route
default via 62.143.72.1 dev eth0.2 proto static src 62.143.72.48
10.0.0.0/8 dev vti0 proto static scope link
62.143.72.0/21 dev eth0.2 proto kernel scope link src 62.143.72.48
192.168.1.0/24 dev br-lan proto kernel scope link src 192.168.1.1
```

```
root@OpenWrt:~# route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
0.0.0.0 62.143.72.1 0.0.0.0 UG 0 0 0 eth0.2
10.0.0.0 0.0.0.0 255.0.0.0 U 0 0 0 vti0
62.143.72.0 0.0.0.0 255.255.248.0 U 0 0 0 eth0.2
192.168.1.0 0.0.0.0 255.255.255.0 U 0 0 0 br-lan
```

```
root@OpenWrt:~# ip route list table 220
root@OpenWrt:~#
```

I don't know how to set up routing. Can you help?

Best Regards

Bernd

#17 - 28.07.2020 09:50 - Tobias Brunner

What's the VTI device for? You could just NAT whatever traffic you want to tunnel to the virtual IP (similar to the rule you install in the updown script), no other traffic will match the IPsec policy and won't get tunneled.

#18 - 28.07.2020 17:49 - Bernd Bernikov

What's the VTI device for?

VTI device is for an OpenWRT interface. It is then possible to create a firewall rule for an interface. With this interface I can block everything in OpenWRT that does not go through the IPsec tunnel.

Best Regards

Bernd.

#19 - 16.08.2020 12:58 - Bernd Bernikov

Can anyone else help?

I don't know how to set manual routing.

When I try to manually set route for table 220, for example

82.199.134.161 is the gateway IP of VPN provider.

10.0.91.66 is the virtual IP of the VPN provider.

```
ip addr add 82.199.134.161 dev vti0
ip addr add 10.0.91.66 dev vti0
ip route add default via 82.199.134.161 dev vti0 proto static src 10.0.91.66 table 220
ip route add 192.168.1.0/24 dev br-lan proto static src 192.168.1.1 table 220
```

Then I get errors.

```
vti0      Link encap:UNSPEC HWaddr C0-A8-01-01-00-00-00-00-00-00-00-00-00-00-00
inet addr:10.0.0.0 P-t-P:10.0.0.0 Mask:255.0.0.0
inet6 addr: fe80::5efe:c0a8:101/64 Scope:Link
UP POINTOPOINT RUNNING NOARP MTU:1416 Metric:1
RX packets:90 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:134 dropped:0 overruns:0 carrier:134
collisions:0 txqueuelen:1000
RX bytes:24010 (23.4 KiB) TX bytes:0 (0.0 B)
```