

VPN over IPsec

Table of Contents

1. Configuring a VPN on FreeBSD 2

Internet Protocol Security (IPsec) is a set of protocols which sit on top of the Internet Protocol (IP) layer. It allows two or more hosts to communicate in a secure manner by authenticating and encrypting each IP packet of a communication session. The FreeBSD IPsec network stack is based on the <http://www.kame.net/> implementation and supports both IPv4 and IPv6 sessions.

IPsec is comprised of the following sub-protocols:

- *Encapsulated Security Payload (ESP)*: this protocol protects the IP packet data from third party interference by encrypting the contents using symmetric cryptography algorithms such as Blowfish and 3DES.
- *Authentication Header (AH)*: this protocol protects the IP packet header from third party interference and spoofing by computing a cryptographic checksum and hashing the IP packet header fields with a secure hashing function. This is then followed by an additional header that contains the hash, to allow the information in the packet to be authenticated.
- *IP Payload Compression Protocol (IPComp)*: this protocol tries to increase communication performance by compressing the IP payload in order to reduce the amount of data sent.


These protocols can either be used together or separately, depending on the environment.

IPsec supports two modes of operation. The first mode, *Transport Mode*, protects communications between two hosts. The second mode, *Tunnel Mode*, is used to build virtual tunnels, commonly known as Virtual Private Networks (VPNs). Consult [ipsec\(4\)](#) for detailed information on the IPsec subsystem in FreeBSD.

The article demonstrates the process of setting up an IPsecVPN between a home network and a corporate network.

In the example scenario:

- Both sites are connected to the Internet through a gateway that is running FreeBSD.
- The gateway on each network has at least one external IP address. In this example, the corporate LAN's external IP address is **172.16.5.4** and the home LAN's external IP address is **192.168.1.12**.
- The internal addresses of the two networks can be either public or private IP addresses. However, the address space must not overlap. In this example, the corporate LAN's internal IP address is **10.246.38.1** and the home LAN's internal IP address is **10.0.0.5**.



```
10.246.38.1/24 -- 172.16.5.4 <--> 192.168.1.12 -- 10.0.0.5/24
```

1. Configuring a VPN on FreeBSD

To begin, [security/ipsec-tools](#) must be installed from the Ports Collection. This software provides a number of applications which support the configuration.

The next requirement is to create two [gif\(4\)](#) pseudo-devices which will be used to tunnel packets and allow both networks to communicate properly. As `root`, run the following command on each gateway:

```
corp-gw# ifconfig gif0 create
corp-gw# ifconfig gif0 10.246.38.1 10.0.0.5
corp-gw# ifconfig gif0 tunnel 172.16.5.4 192.168.1.12
```

```
home-gw# ifconfig gif0 create
home-gw# ifconfig gif0 10.0.0.5 10.246.38.1
home-gw# ifconfig gif0 tunnel 192.168.1.12 172.16.5.4
```

Verify the setup on each gateway, using `ifconfig gif0`. Here is the output from the home gateway:

```
gif0: flags=8051 mtu 1280
tunnel inet 172.16.5.4 --> 192.168.1.12
inet6 fe80::2e0:81ff:fe02:5881%gif0 prefixlen 64 scopeid 0x6
inet 10.246.38.1 --> 10.0.0.5 netmask 0xffffffff00
```

Here is the output from the corporate gateway:

```
gif0: flags=8051 mtu 1280
tunnel inet 192.168.1.12 --> 172.16.5.4
inet 10.0.0.5 --> 10.246.38.1 netmask 0xffffffff00
inet6 fe80::250:bfff:fe3a:c1f%gif0 prefixlen 64 scopeid 0x4
```

Once complete, both internal IP addresses should be reachable using [ping\(8\)](#):

```
home-gw# ping 10.0.0.5
PING 10.0.0.5 (10.0.0.5): 56 data bytes
64 bytes from 10.0.0.5: icmp_seq=0 ttl=64 time=42.786 ms
64 bytes from 10.0.0.5: icmp_seq=1 ttl=64 time=19.255 ms
64 bytes from 10.0.0.5: icmp_seq=2 ttl=64 time=20.440 ms
64 bytes from 10.0.0.5: icmp_seq=3 ttl=64 time=21.036 ms
--- 10.0.0.5 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 19.255/25.879/42.786/9.782 ms
```

```
corp-gw# ping 10.246.38.1
PING 10.246.38.1 (10.246.38.1): 56 data bytes
64 bytes from 10.246.38.1: icmp_seq=0 ttl=64 time=28.106 ms
64 bytes from 10.246.38.1: icmp_seq=1 ttl=64 time=42.917 ms
64 bytes from 10.246.38.1: icmp_seq=2 ttl=64 time=127.525 ms
64 bytes from 10.246.38.1: icmp_seq=3 ttl=64 time=119.896 ms
64 bytes from 10.246.38.1: icmp_seq=4 ttl=64 time=154.524 ms
--- 10.246.38.1 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 28.106/94.594/154.524/49.814 ms
```

As expected, both sides have the ability to send and receive ICMP packets from the privately configured addresses. Next, both gateways must be told how to route packets in order to correctly send traffic from the networks behind each gateway. The following commands will achieve this goal:

```
corp-gw# route add 10.0.0.0 10.0.0.5 255.255.255.0
corp-gw# route add net 10.0.0.0: gateway 10.0.0.5
home-gw# route add 10.246.38.0 10.246.38.1 255.255.255.0
home-gw# route add host 10.246.38.0: gateway 10.246.38.1
```

Internal machines should be reachable from each gateway as well as from machines behind the gateways. Again, use [ping\(8\)](#) to confirm:

```
corp-gw# ping -c 3 10.0.0.8
PING 10.0.0.8 (10.0.0.8): 56 data bytes
64 bytes from 10.0.0.8: icmp_seq=0 ttl=63 time=92.391 ms
64 bytes from 10.0.0.8: icmp_seq=1 ttl=63 time=21.870 ms
64 bytes from 10.0.0.8: icmp_seq=2 ttl=63 time=198.022 ms
--- 10.0.0.8 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 21.870/101.846/198.022/74.001 ms

home-gw# ping -c 3 10.246.38.107
PING 10.246.38.1 (10.246.38.107): 56 data bytes
64 bytes from 10.246.38.107: icmp_seq=0 ttl=64 time=53.491 ms
64 bytes from 10.246.38.107: icmp_seq=1 ttl=64 time=23.395 ms
64 bytes from 10.246.38.107: icmp_seq=2 ttl=64 time=23.865 ms
--- 10.246.38.107 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 21.145/31.721/53.491/12.179 ms
```

At this point, traffic is flowing between the networks encapsulated in a gif tunnel but without any encryption. Next, use IPsec to encrypt traffic using pre-shared keys (PSK). Other than the IP addresses, `/usr/local/etc/racoon/racoon.conf` on both gateways will be identical and look similar to:

```

path    pre_shared_key  "/usr/local/etc/racoon/psk.txt"; #location of pre-shared key
file
log     debug; #log verbosity setting: set to 'notify' when testing and debugging is
complete

padding # options are not to be changed
{
    maximum_length  20;
    randomize        off;
    strict_check     off;
    exclusive_tail   off;
}

timer  # timing options. change as needed
{
    counter          5;
    interval         20 sec;
    persend          1;
#    natt_keepalive  15 sec;
    phase1           30 sec;
    phase2           15 sec;
}

listen # address [port] that racoon will listen on
{
    isakmp           172.16.5.4 [500];
    isakmp_natt      172.16.5.4 [4500];
}

remote 192.168.1.12 [500]
{
    exchange_mode    main,aggressive;
    doi              ipsec_doi;
    situation         identity_only;
    my_identifier     address 172.16.5.4;
    peers_identifier  address 192.168.1.12;
    lifetime          time 8 hour;
    passive           off;
    proposal_check    obey;
#    nat_traversal    off;
    generate_policy   off;

                    proposal {
                        encryption_algorithm  blowfish;
                        hash_algorithm        md5;
                        authentication_method  pre_shared_key;
                        lifetime time         30 sec;
                        dh_group              1;
                    }
}

```

```
sainfo (address 10.246.38.0/24 any address 10.0.0.0/24 any) # address
$network/$netmask $type address $network/$netmask $type ( $type being any or esp)
{
    # $network must be the two internal networks you are
    joining.
    pfs_group      1;
    lifetime       time    36000 sec;
    encryption_algorithm    blowfish,3des;
    authentication_algorithm    hmac_md5,hmac_sha1;
    compression_algorithm    deflate;
}
```

For descriptions of each available option, refer to the manual page for `racoon.conf`.

The Security Policy Database (SPD) needs to be configured so that FreeBSD and `racoon` are able to encrypt and decrypt network traffic between the hosts.

This can be achieved with a shell script, similar to the following, on the corporate gateway. This file will be used during system initialization and should be saved as `/usr/local/etc/racoon/setkey.conf`.

```
flush;
spdflush;
# To the home network
spdadd 10.246.38.0/24 10.0.0.0/24 any -P out ipsec esp/tunnel/172.16.5.4-
192.168.1.12/use;
spdadd 10.0.0.0/24 10.246.38.0/24 any -P in ipsec esp/tunnel/192.168.1.12-
172.16.5.4/use;
```

Once in place, `racoon` may be started on both gateways using the following command:

```
# /usr/local/sbin/racoon -F -f /usr/local/etc/racoon/racoon.conf -l
/var/log/racoon.log
```

The output should be similar to the following:

```
corp-gw# /usr/local/sbin/racoon -F -f /usr/local/etc/racoon/racoon.conf
Foreground mode.
2006-01-30 01:35:47: INFO: begin Identity Protection mode.
2006-01-30 01:35:48: INFO: received Vendor ID: KAME/racoon
2006-01-30 01:35:55: INFO: received Vendor ID: KAME/racoon
2006-01-30 01:36:04: INFO: ISAKMP-SA established 172.16.5.4[500]-192.168.1.12[500]
spi:623b9b3bd2492452:7deab82d54ff704a
2006-01-30 01:36:05: INFO: initiate new phase 2 negotiation:
172.16.5.4[0]192.168.1.12[0]
2006-01-30 01:36:09: INFO: IPsec-SA established: ESP/Tunnel 192.168.1.12[0]-
>172.16.5.4[0] spi=28496098(0x1b2d0e2)
2006-01-30 01:36:09: INFO: IPsec-SA established: ESP/Tunnel 172.16.5.4[0]-
>192.168.1.12[0] spi=47784998(0x2d92426)
```

```
2006-01-30 01:36:13: INFO: respond new phase 2 negotiation:
172.16.5.4[0]192.168.1.12[0]
2006-01-30 01:36:18: INFO: IPsec-SA established: ESP/Tunnel 192.168.1.12[0]-
>172.16.5.4[0] spi=124397467(0x76a279b)
2006-01-30 01:36:18: INFO: IPsec-SA established: ESP/Tunnel 172.16.5.4[0]-
>192.168.1.12[0] spi=175852902(0xa7b4d66)
```

To ensure the tunnel is working properly, switch to another console and use `tcpdump(1)` to view network traffic using the following command. Replace `em0` with the network interface card as required:

```
corp-gw# tcpdump -i em0 host 172.16.5.4 and dst 192.168.1.12
```

Data similar to the following should appear on the console. If not, there is an issue and debugging the returned data will be required.

```
01:47:32.021683 IP corporatenetwork.com > 192.168.1.12.privatenetwork.com:
ESP(spi=0x02acbf9f,seq=0xa)
01:47:33.022442 IP corporatenetwork.com > 192.168.1.12.privatenetwork.com:
ESP(spi=0x02acbf9f,seq=0xb)
01:47:34.024218 IP corporatenetwork.com > 192.168.1.12.privatenetwork.com:
ESP(spi=0x02acbf9f,seq=0xc)
```

At this point, both networks should be available and seem to be part of the same network. Most likely both networks are protected by a firewall. To allow traffic to flow between them, rules need to be added to pass packets. For the `ipfw(8)` firewall, add the following lines to the firewall configuration file:

```
ipfw add 00201 allow log esp from any to any
ipfw add 00202 allow log ah from any to any
ipfw add 00203 allow log ipencap from any to any
ipfw add 00204 allow log udp from any 500 to any
```



The rule numbers may need to be altered depending on the current host configuration.

For users of `pf(4)` or `ipf(8)`, the following rules should do the trick:

```
pass in quick proto esp from any to any
pass in quick proto ah from any to any
pass in quick proto ipencap from any to any
pass in quick proto udp from any port = 500 to any port = 500
pass in quick on gif0 from any to any
pass out quick proto esp from any to any
pass out quick proto ah from any to any
```

```
pass out quick proto ipencap from any to any
pass out quick proto udp from any port = 500 to any port = 500
pass out quick on gif0 from any to any
```

Finally, to allow the machine to start support for the VPN during system initialization, add the following lines to `/etc/rc.conf`:

```
ipsec_enable="YES"
ipsec_program="/usr/local/sbin/setkey"
ipsec_file="/usr/local/etc/racoon/setkey.conf" # allows setting up spd policies on
boot
racoon_enable="yes"
```