LF NETWORKING
Developer & Testing Forum

IPSec Acceleration with VPP-Sswan and Linux-CP

Kai Ji - Intel

Gabriel Oginski - Mobica

https://lfnetworking.org

Introduction

VPP IPsec in FD.io offers secure and fast networking applications. And easy-to-use commands configure SPD, SA, and cryptographic settings in VPP.

1.89 Terabit NDR on 4th Gen Intel Xeon with Intel® Multi-Buffer Crypto

- vAES and vPCLMUL instruction
- AVX512 allows four 128-bits AES blocks in parallel
- Up to 50Gbps for a single SA with tunnel-based IPsec in AES-GCM-128

FD.io VPP IPsec has an incomplete IKEv2 implementation and relies on StrongSwan. Offloading to VPP IPsec provides significant performance benefits.

Connect Kernel and User Space I/O AF_PACKET or AF_XDP

- StrongSwan must be updated to understand VPP af_packet interfaces.
- A single routing entry needs to be configured for both kernel and VPP separately.
- Even if everything is working properly, the overall IPsec throughput will be limited by slow af_packet interfaces.

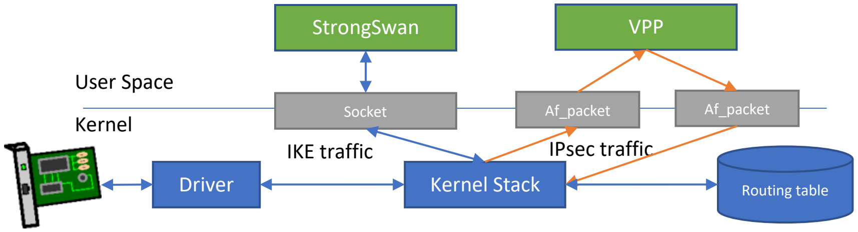

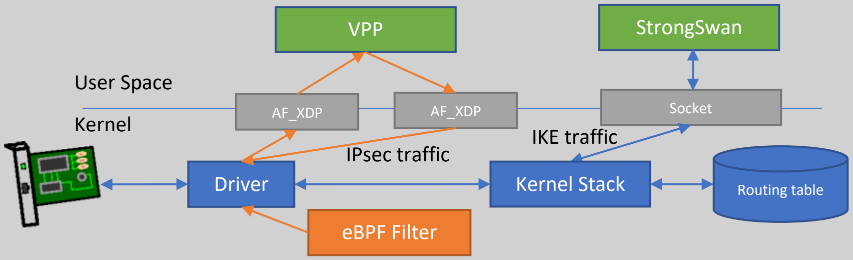

- AF_XDP requires Linux kernel (4.18) and network adapter kernel driver support.
- An eBPF filter program is static and wont updated during running.

VPP-SSwan with Linux Control Plane

VPP-SSwan plug-in

- Written to following the StrongSwan Charon spec
- Network I/O and IPsec data path are processed by VPP

Linux Control Plane plugin

- Traffic is cross connected between the Linux tun/tap interface and NIC
- ARP/ND table updated automatically

The VPP-SSwan plugin is included in VPP from 22.10 release and works with StrongSwan 5.9.5 onwards.

VPP-SSwan with Linux Control Plane

LF NETWORKING
LFN Developer & Testing Forum

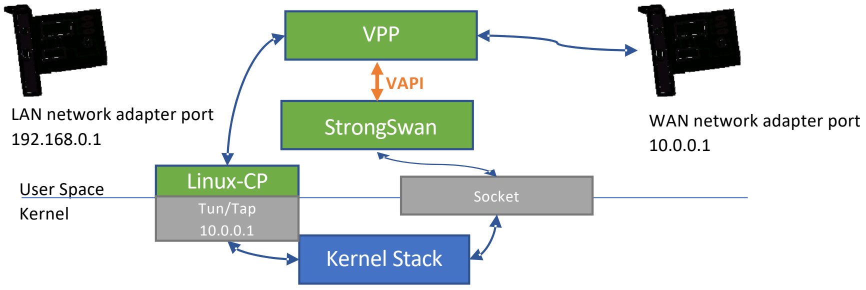

IKE Traffic Between VPP Owned Network Adapter and StrongSwan Application

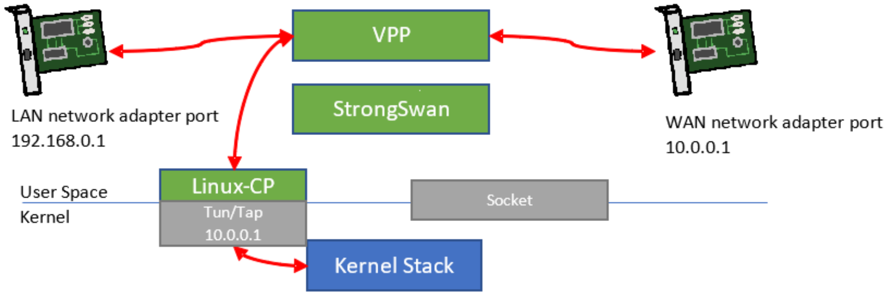

IPsec Traffic Between VPP Owned Network Adapter and Kernel Owned Tun/Tap Port

Demo

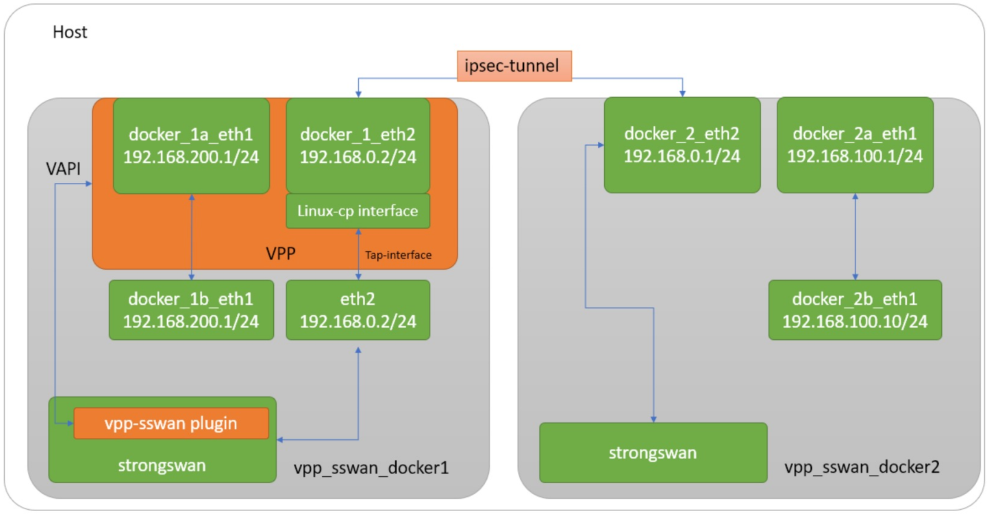

```
path: vpp/extras/strongswan/vpp_sswan/docker
```

```
./run.sh prepare_containers
```

(to downloaded image of ubuntu, prepared image docker and created containers)

```
./run.sh config
```

(to config virtual pairs ethernet NIC's in kernel – we don’t need physical NIC’s, set ip address, routed and etc., run VPP and StrongSwan, checked connection between them, initialized and established connection by StrongSwan)

```
./run.sh clean
```

(to clean-up after executed ./run.sh config
– terminated connection between peer, stopped VPP and clean-up virtual interfaces)

```
./run.sh deleted
```

(to clean-up after executed ./run.sh prepare_containers
– clean-up each containers, deleted docker image)

LFN Developer & Testing Forum

Q&A

LF NETWORKING

LFN Developer & Testing Forum