Encrypting OVN tunnels with IPsec

Qiuyu Xiao (qiuyu.xiao.qyx@gmail.com)

Ben Pfaff (blp@ovn.org)

Open Virtual Network (OVN)

OVN provides a logical network abstraction on top of a physical network

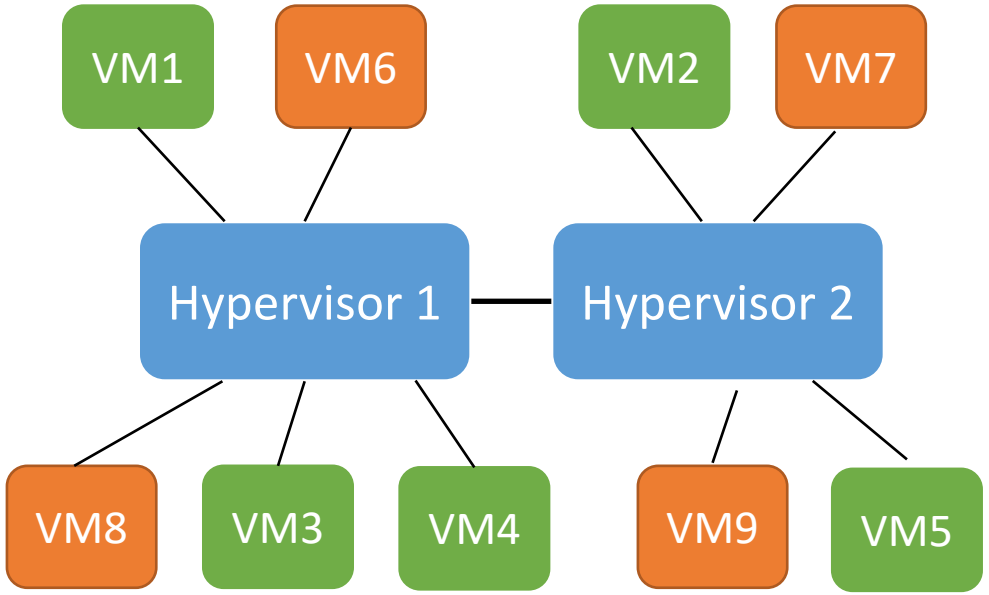

Physical

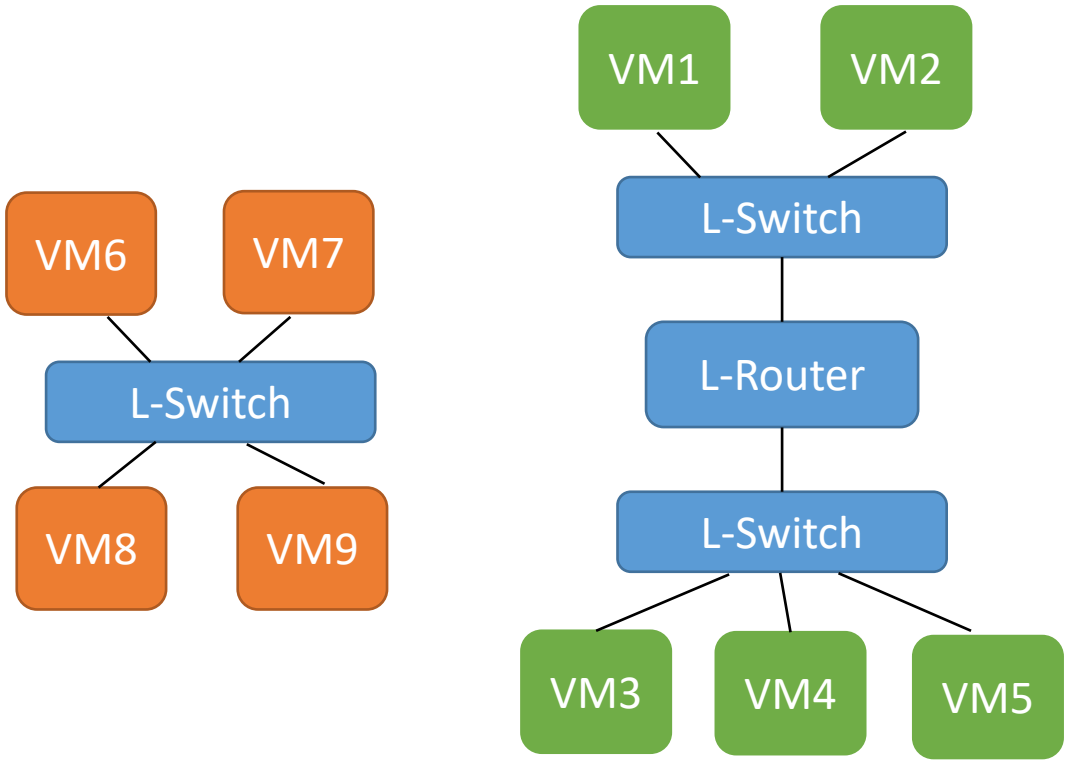

Logical

Open Virtual Network (OVN)

VMs are oblivious to the physical network states

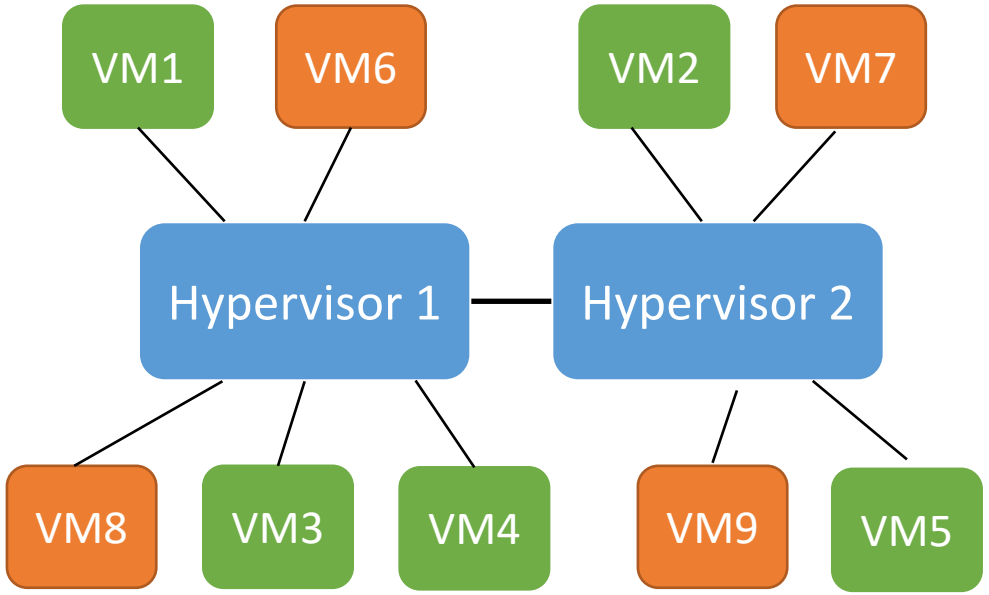

Physical

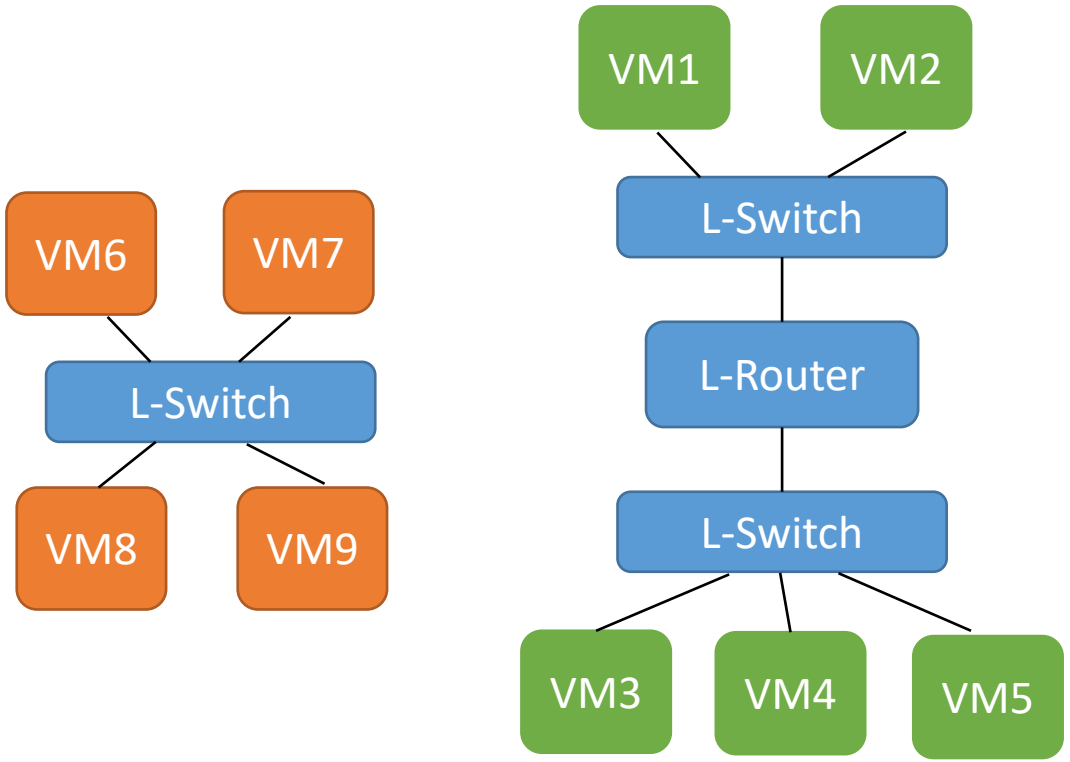

Logical

Open Virtual Network (OVN)

Network appliances can be implemented and placed in the logical network

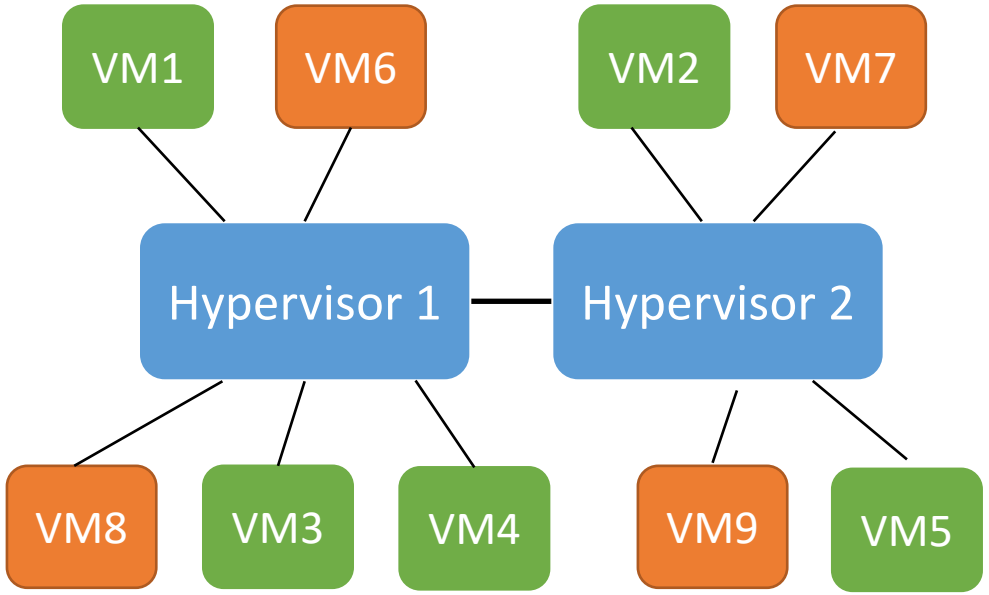

Physical

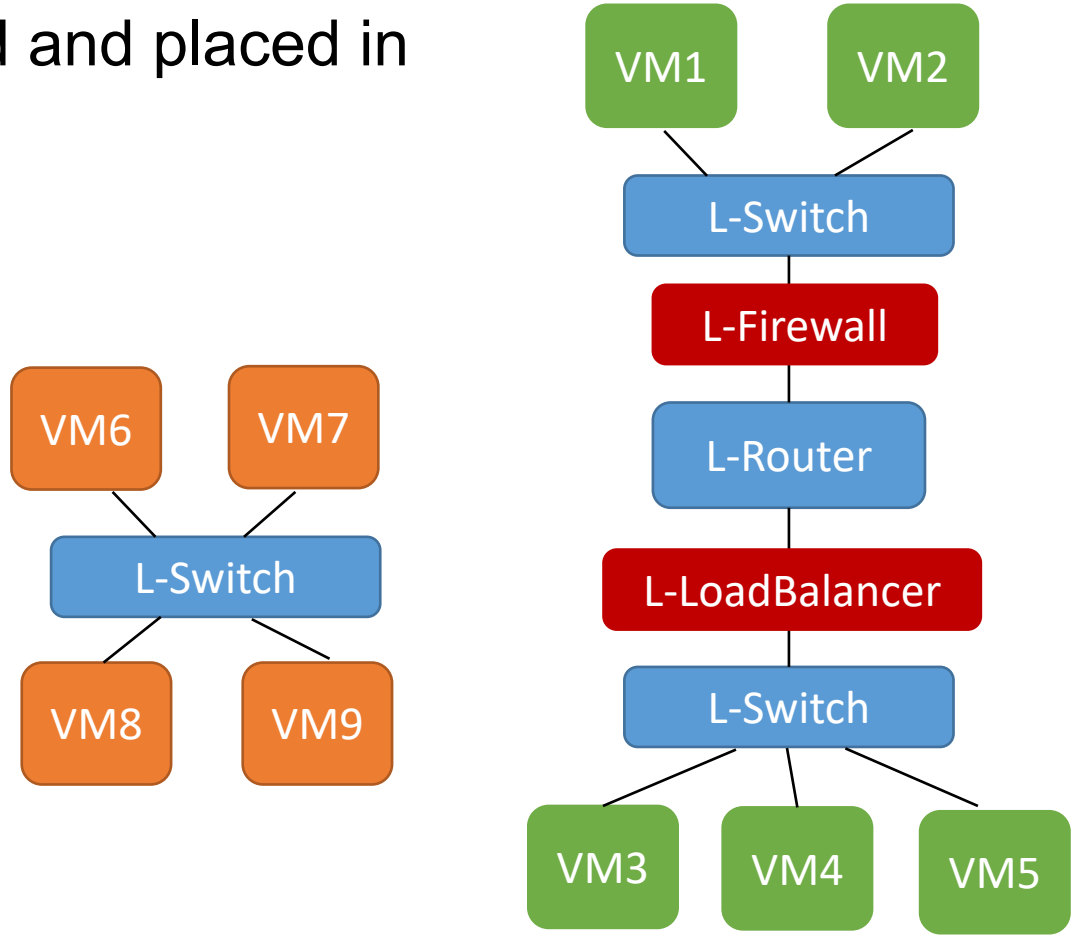

Logical

OVN Tunnel Traffic

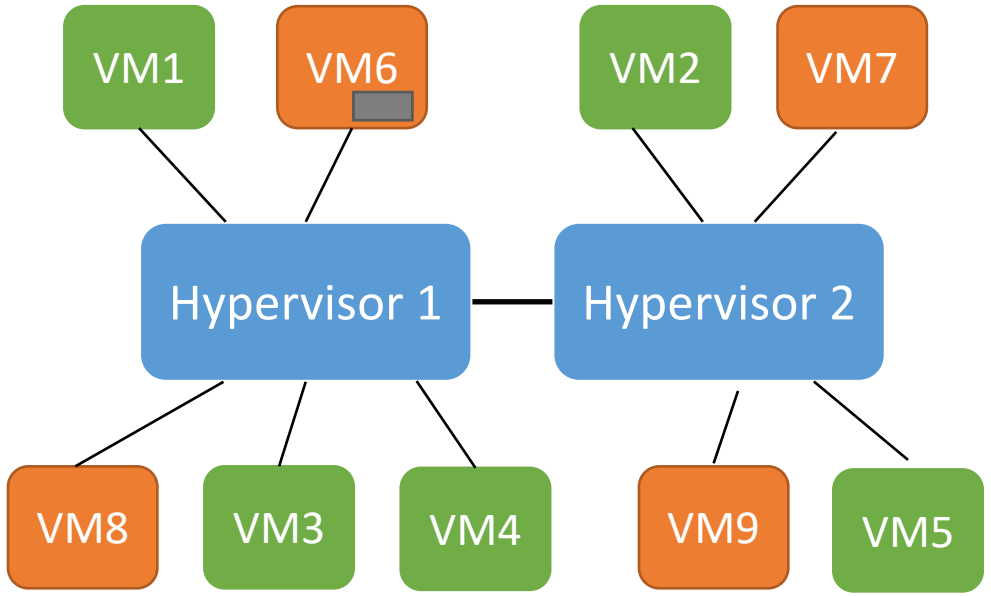

OVN Tunnel Traffic

Outer Ethernet Header	Outer IP Header	Outer UDP Header	Geneve Header	Inner Ethernet Header	Inner IP Header	Payload

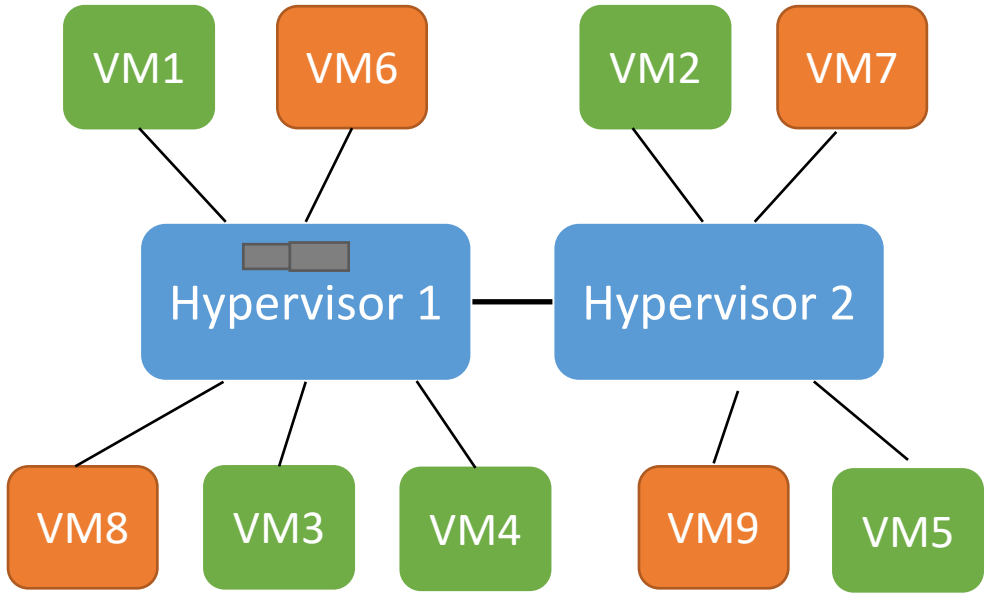

OVN Tunnel Traffic

Outer Ethernet Header	Outer IP Header	Outer UDP Header	Geneve Header	Inner Ethernet Header	Inner IP Header	Payload

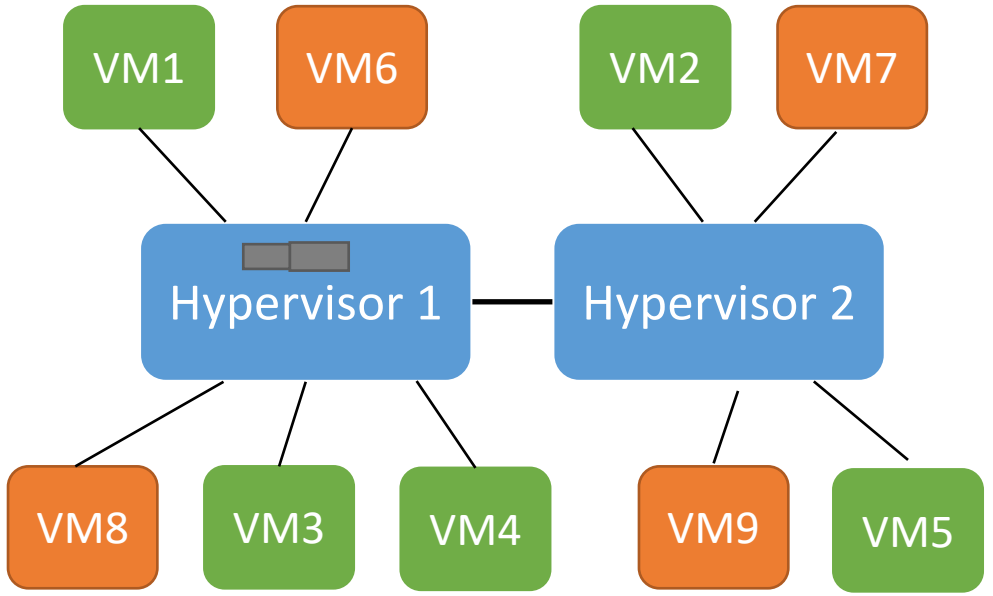

OVN Tunnel Traffic

Outer Ethernet Header	Outer IP Header	Outer UDP Header	Geneve Header	Inner Ethernet Header	Inner IP Header	Payload

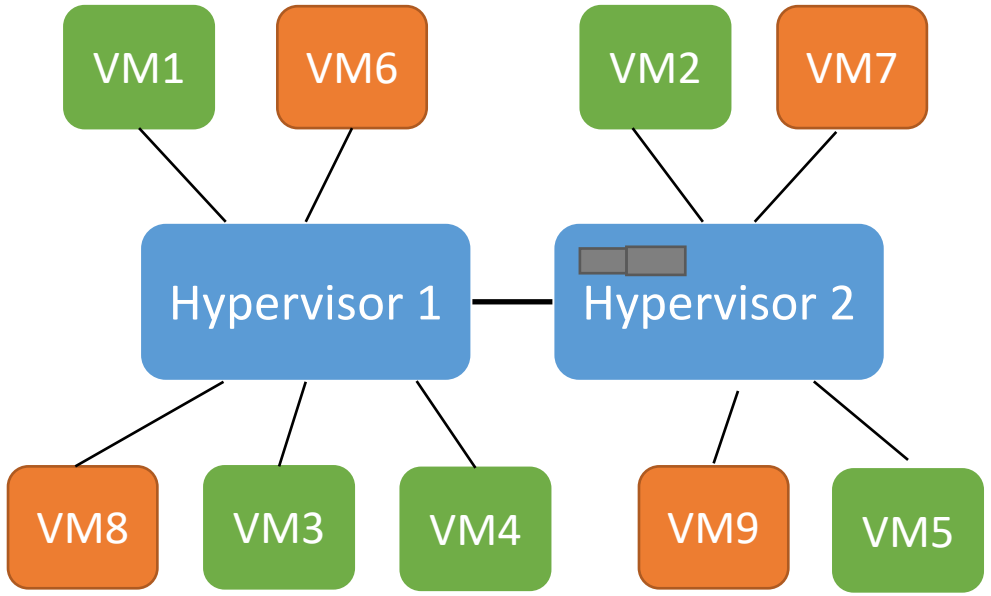

OVN Tunnel Traffic

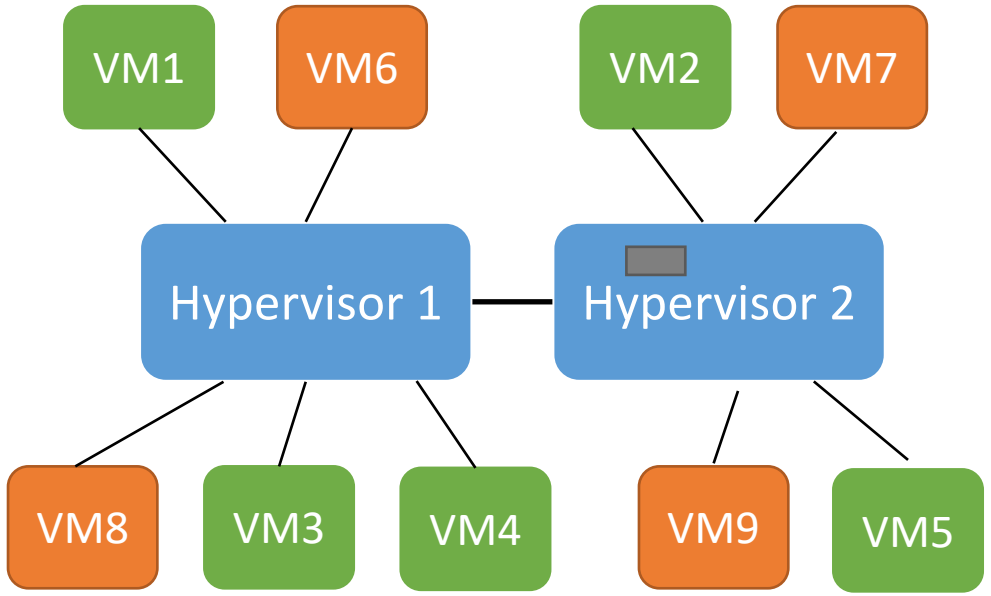

The Needs for Tunnel Encryption

- VMs compute and communicate sensitive data, e.g., financial and health data
- Physical network devices (e.g., router, switch) cannot be trusted or might be compromised
  - ❑ Traffic across datacenters
  - ❑ Router misconfiguration
  - ❑ Attackers breaking into internal network
  - ❑ Phishing or social engineering attacks on administrators

Encrypting Tunnel Traffic with IPsec

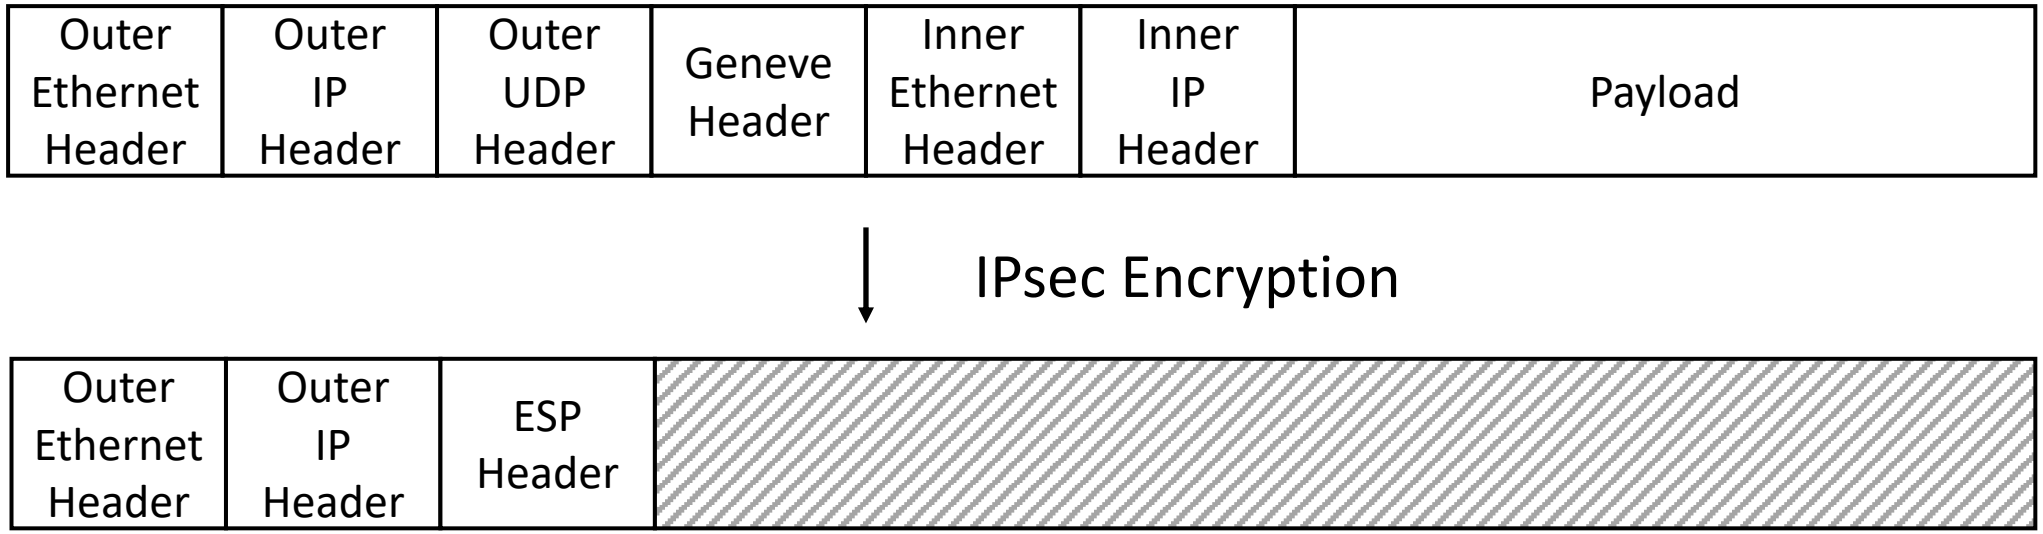

- Confidentiality
- Integrity
- Authenticity

IPsec in Linux

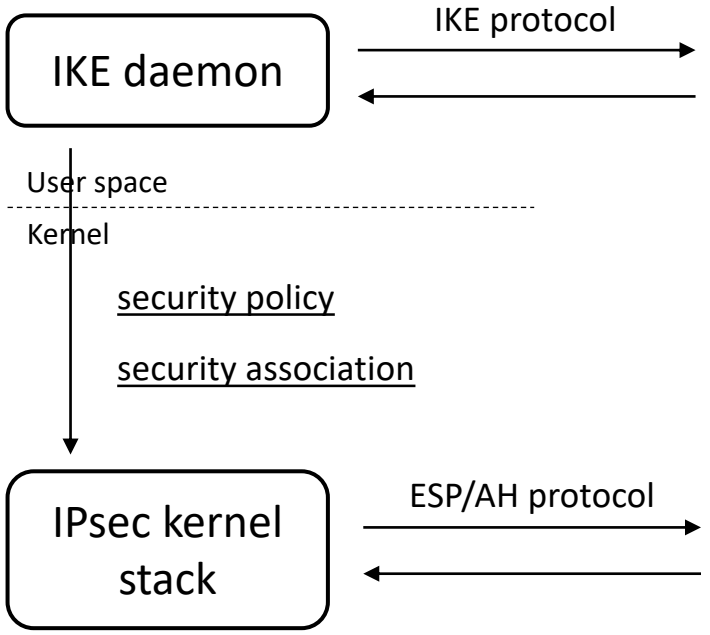

IPsec in Linux

IKE daemon

- Authentication
- Negotiates cryptographic algorithms
- Generates keying material

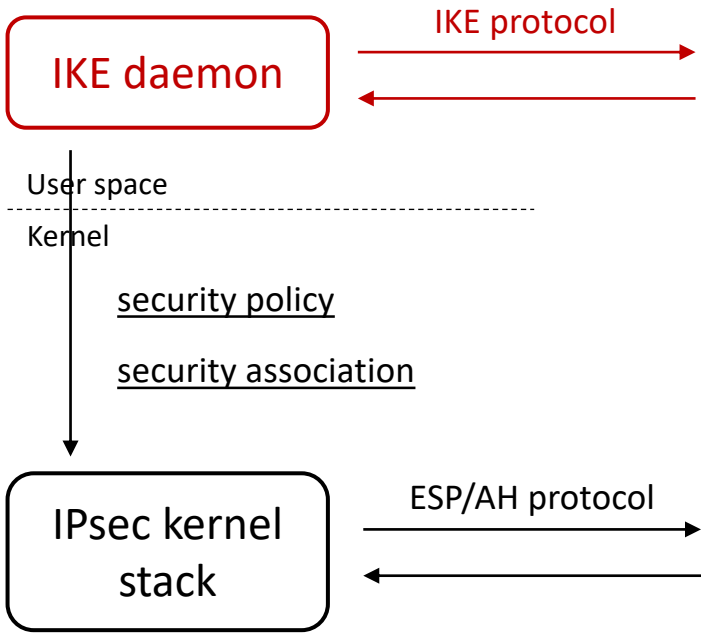

IPsec in Linux

IKE daemon

- Authentication
- Negotiates cryptographic algorithms
- Generates keying material
- Installs security policy and security association

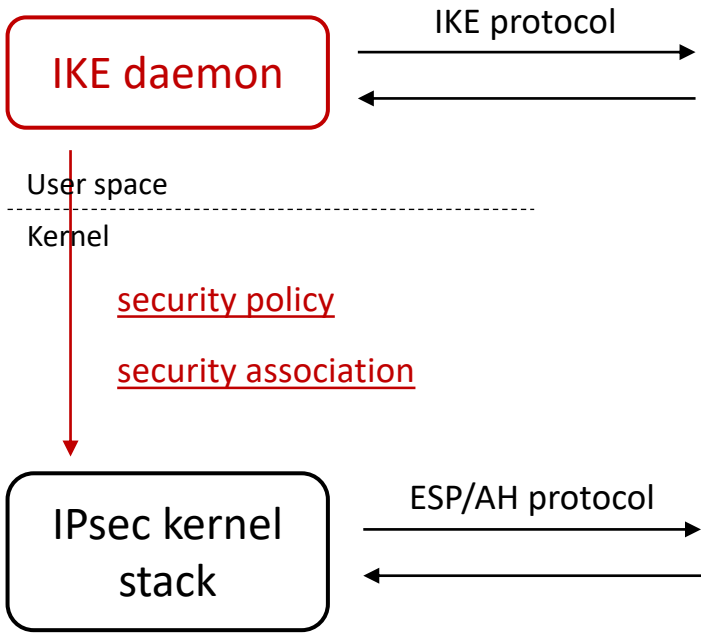

IPsec in Linux

IKE daemon

- Authentication
- Negotiates cryptographic algorithms
- Generates keying material
- Installs security policy and security association

Which traffic to protect

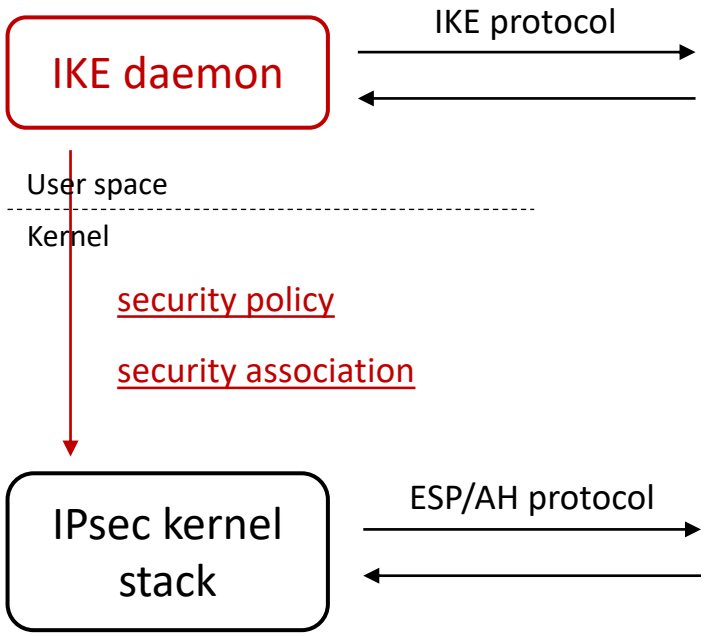

IPsec in Linux

IKE daemon

- Authentication
- Negotiates cryptographic algorithms
- Generates keying material
- Installs security policy and security association

How to protect the selected traffic

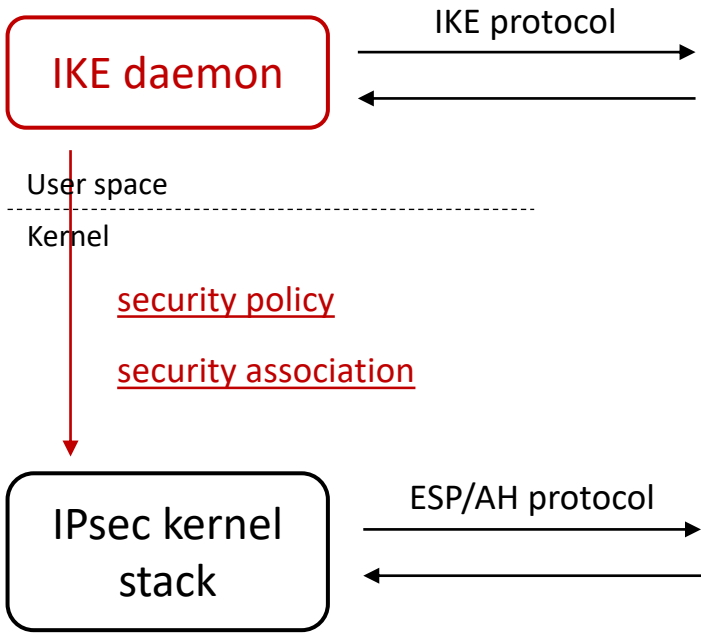

IPsec in Linux

IPsec kernel stack

- Encryption and decryption
- Checks integrity and authenticity

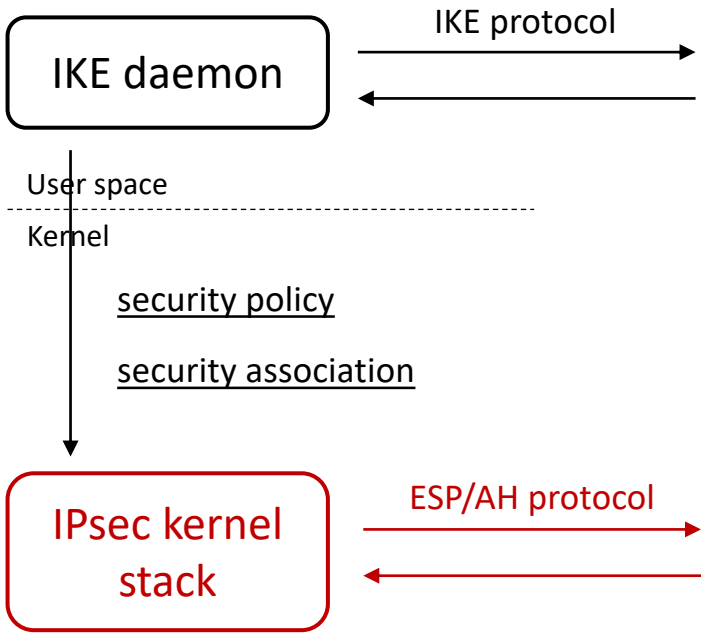

OVS IPsec Tunnel

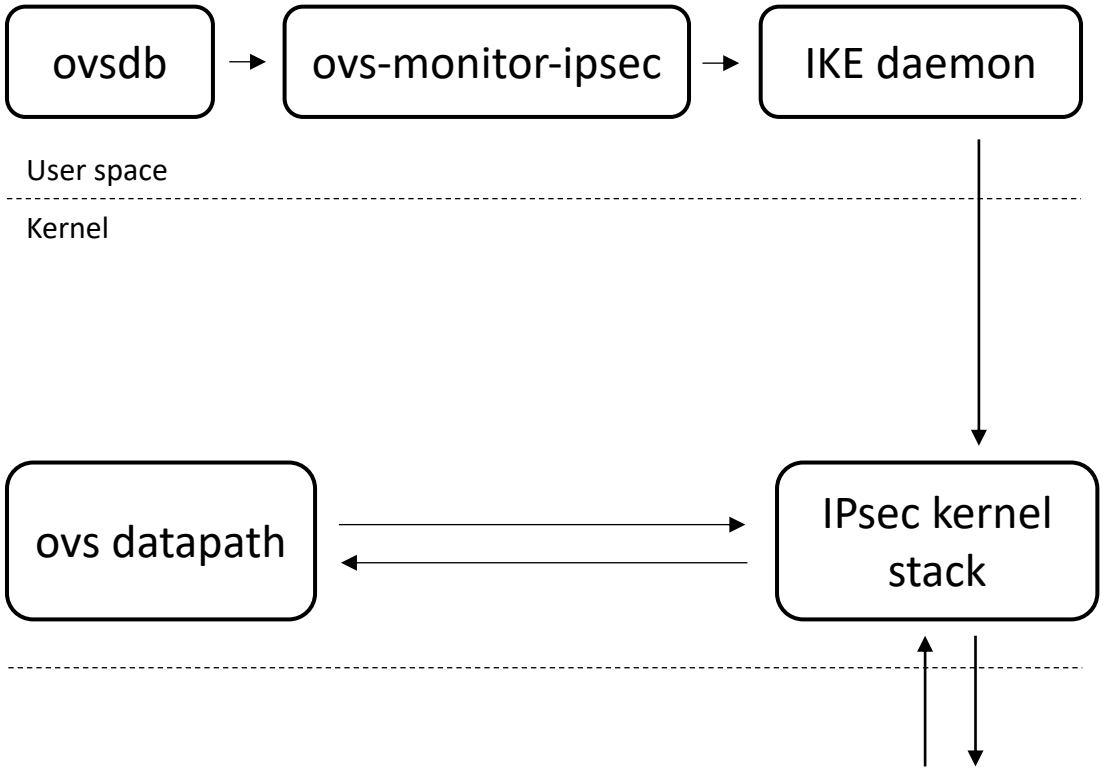

OVS IPsec Tunnel

Configuring IPsec tunnel via ovsdb

- Using pre-shared key

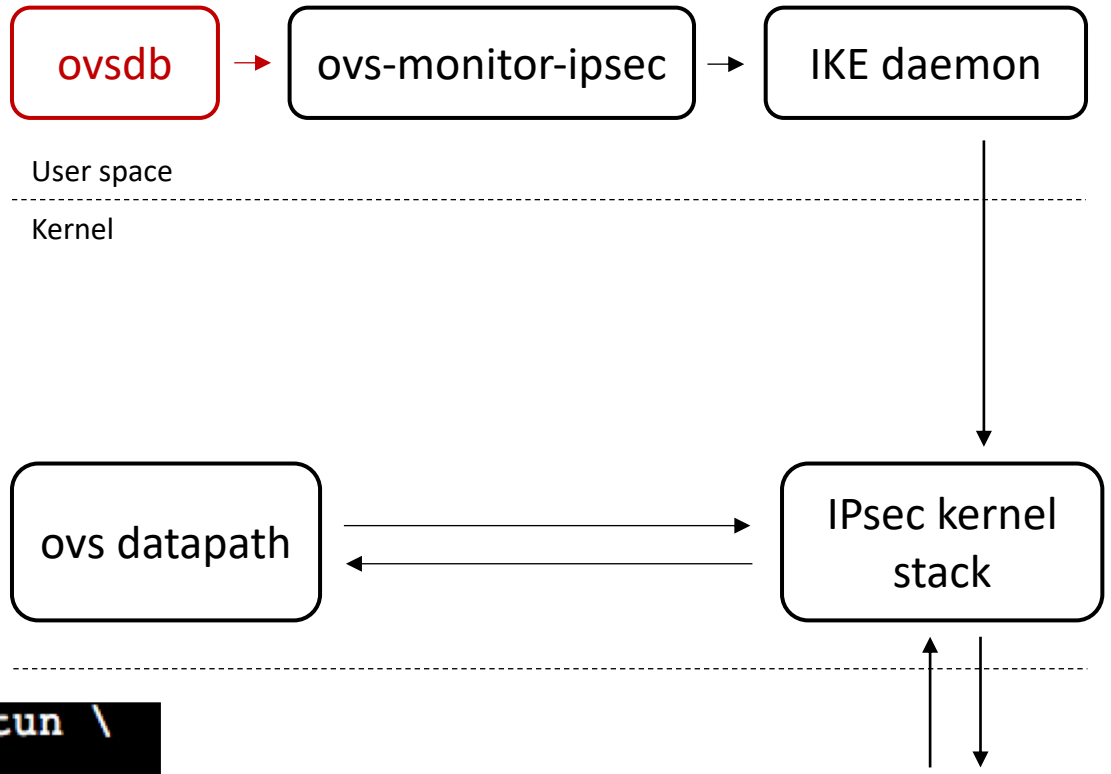

For example:

```
root@ubuntu:~# ovs-vsctl add-port br-int tun \
> -- set interface tun type=geneve \
> options:local_ip=10.33.78.172 \
> options:remote_ip=10.33.79.149 \
> options:psk=swordfish
```

OVS IPsec Tunnel

Configuring IPsec tunnel via ovsdb

- Using pre-shared key
- Using self-signed certificate

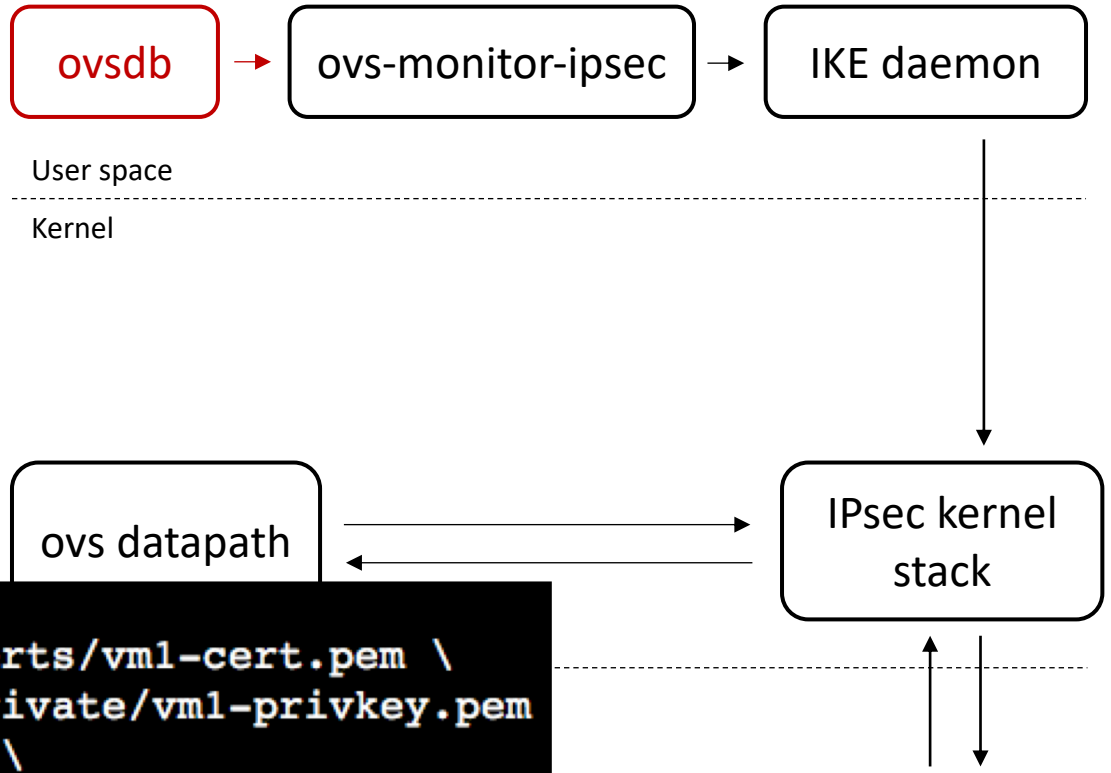

For example:

```
root@vm1:~# ovs-vsctl set Open_vSwitch . \
> other_config:certificate=/etc/ipsec.d/certs/vm1-cert.pem \
> other_config:private_key=/etc/ipsec.d/private/vm1-privkey.pem
root@vm1:~# ovs-vsctl add-port br-int tun \
> -- set interface tun type=geneve \
> options:local_ip=10.33.78.172 \
> options:remote_ip=10.33.79.149 \
> options:remote_cert=/etc/ipsec.d/certs/vm2-cert.pem
```

OVS IPsec Tunnel

Configuring IPsec tunnel via ovsdb

- Using pre-shared key
- Using self-signed certificate
- Using CA-signed certificate

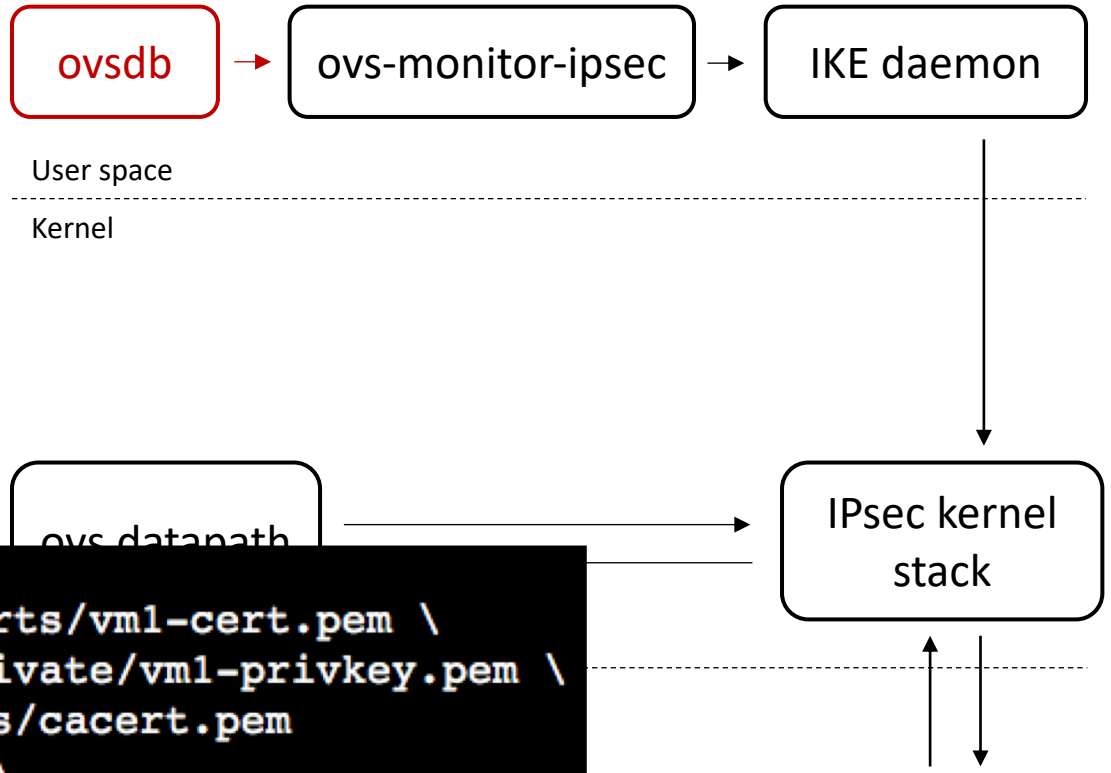

For example:

```
root@vm1:~# ovs-vsctl set Open_vSwitch . \
> other_config:certificate=/etc/ipsec.d/certs/vm1-cert.pem \
> other_config:private_key=/etc/ipsec.d/private/vm1-privkey.pem \
> other_config:ca_cert=/etc/ipsec.d/cacerts/cacert.pem
root@vm1:~# ovs-vsctl add-port br-int tun \
> -- set interface tun type=geneve \
> options:local_ip=10.33.78.172 \
> options:remote_ip=10.33.79.149 \
> options:remote_name=vm2
```

OVS IPsec Tunnel

Establishing IPsec tunnel

- ovs-monitor-ipsec configures IKE daemon

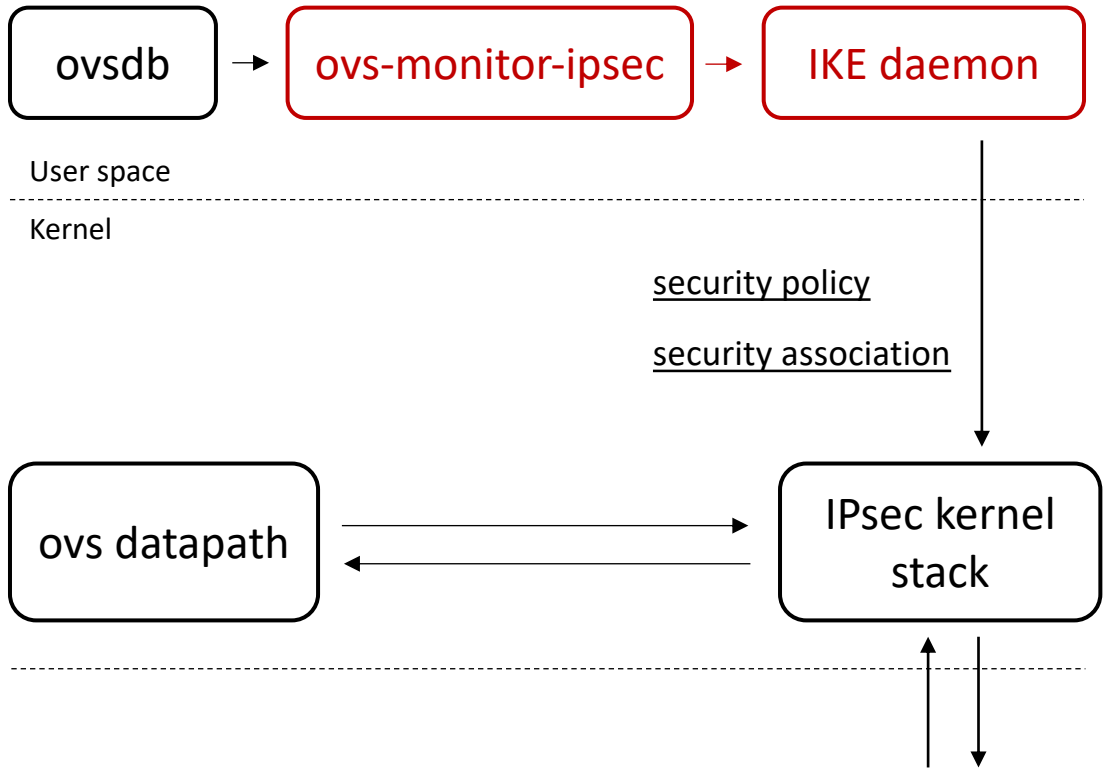

OVS IPsec Tunnel

Establishing IPsec tunnel

- ovs-monitor-ipsec configures IKE daemon
- IKE daemon sets up security policy and security association

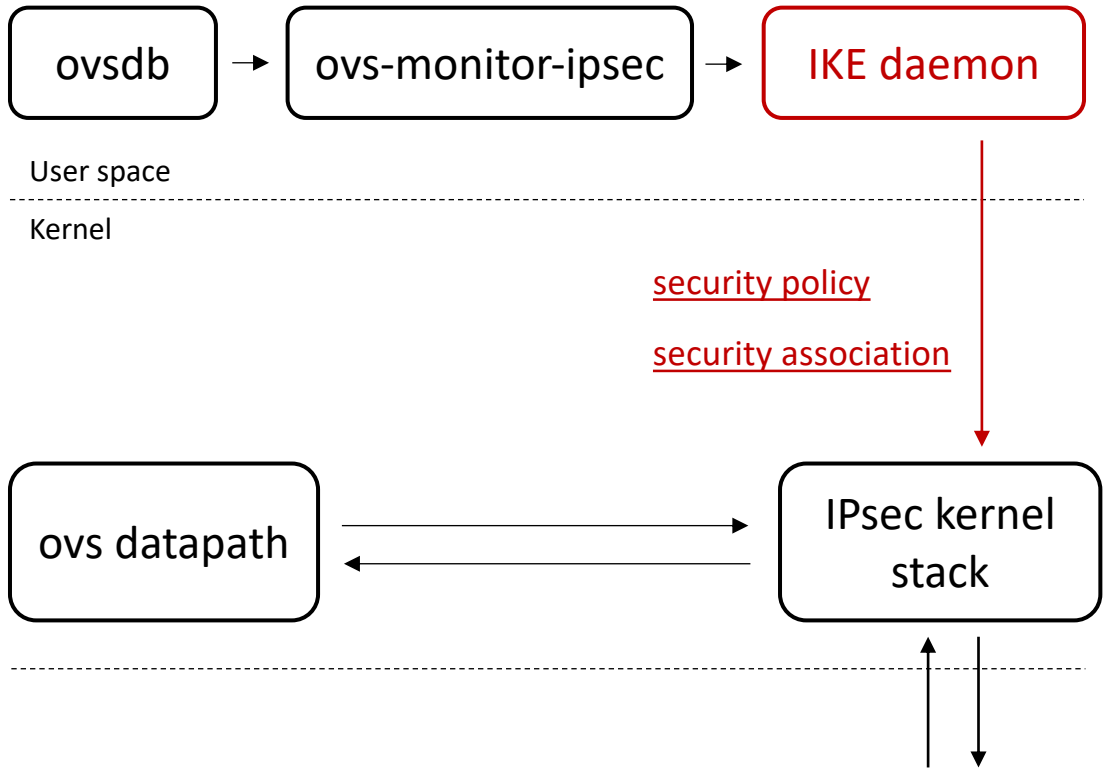

OVS IPsec Tunnel

Establishing IPsec tunnel

- ovs-monitor-ipsec configures IKE daemon
- IKE daemon sets up security policy and security association

For example (geneve tunnel):

```
root@ubuntu:~/debian/4.13# ip xfrm policy show
src 10.33.78.172/32 dst 10.33.79.149/32 proto udp sport 6081
        dir in priority 5888
        tmpl src 0.0.0.0 dst 0.0.0.0
                proto esp reqid 2 mode transport
src 10.33.79.149/32 dst 10.33.78.172/32 proto udp dport 6081
        dir out priority 5888
        tmpl src 0.0.0.0 dst 0.0.0.0
                proto esp reqid 2 mode transport
src 10.33.78.172/32 dst 10.33.79.149/32 proto udp dport 6081
        dir in priority 5888
        tmpl src 0.0.0.0 dst 0.0.0.0
                proto esp reqid 1 mode transport
src 10.33.79.149/32 dst 10.33.78.172/32 proto udp sport 6081
        dir out priority 5888
        tmpl src 0.0.0.0 dst 0.0.0.0
                proto esp reqid 1 mode transport
```

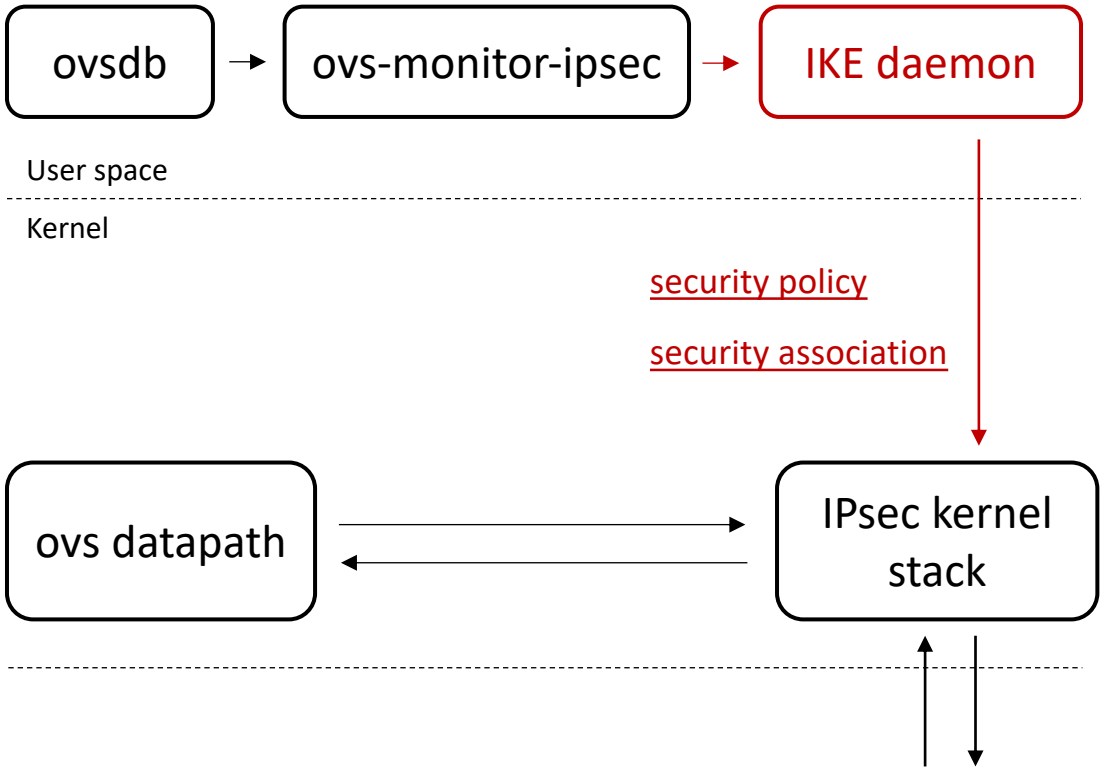

OVS IPsec Tunnel

IPsec kernel stack

- Encryption and decryption
- Checks integrity and authenticity

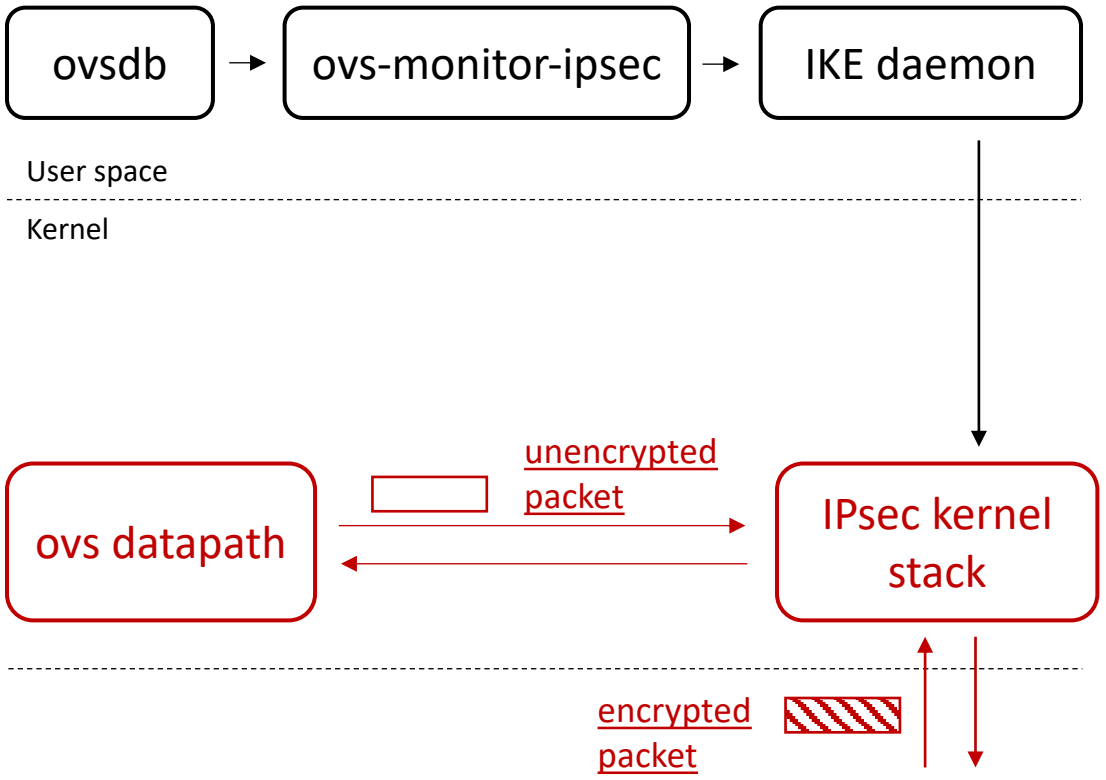

OVN IPsec

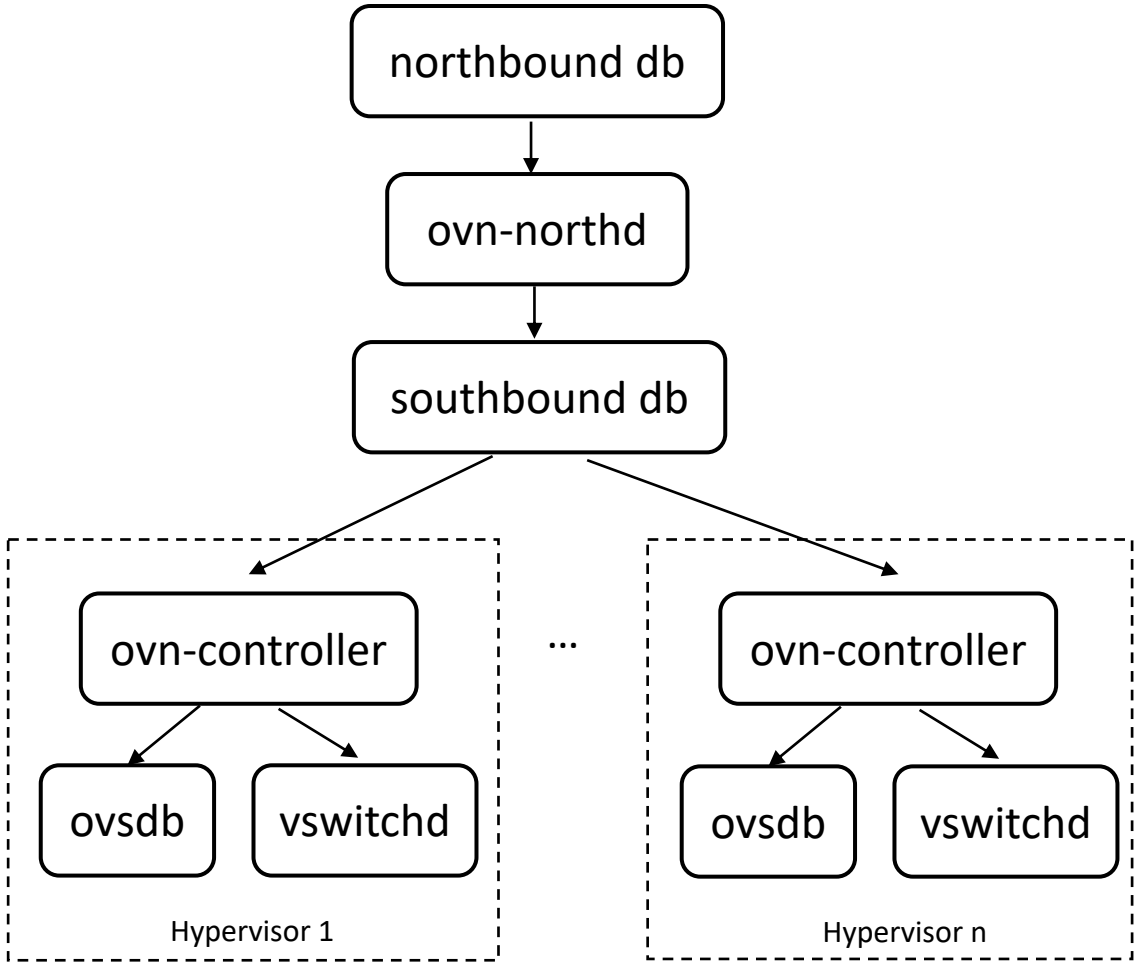
northbound db
ovn-northd
southbound db
ovn-controller
ovsdb
vswitchd
Hypervisor 1
...
ovn-controller
ovsdb
vswitchd
Hypervisor n

OVN IPsec

- In each hypervisor, configure ovsdb to use CA-signed certificate for authentication
- Enable IPsec by configuring northbound database

For example:

```
root@ubuntu:~# ovn-nbctl set nb_global . ipsec=true
```

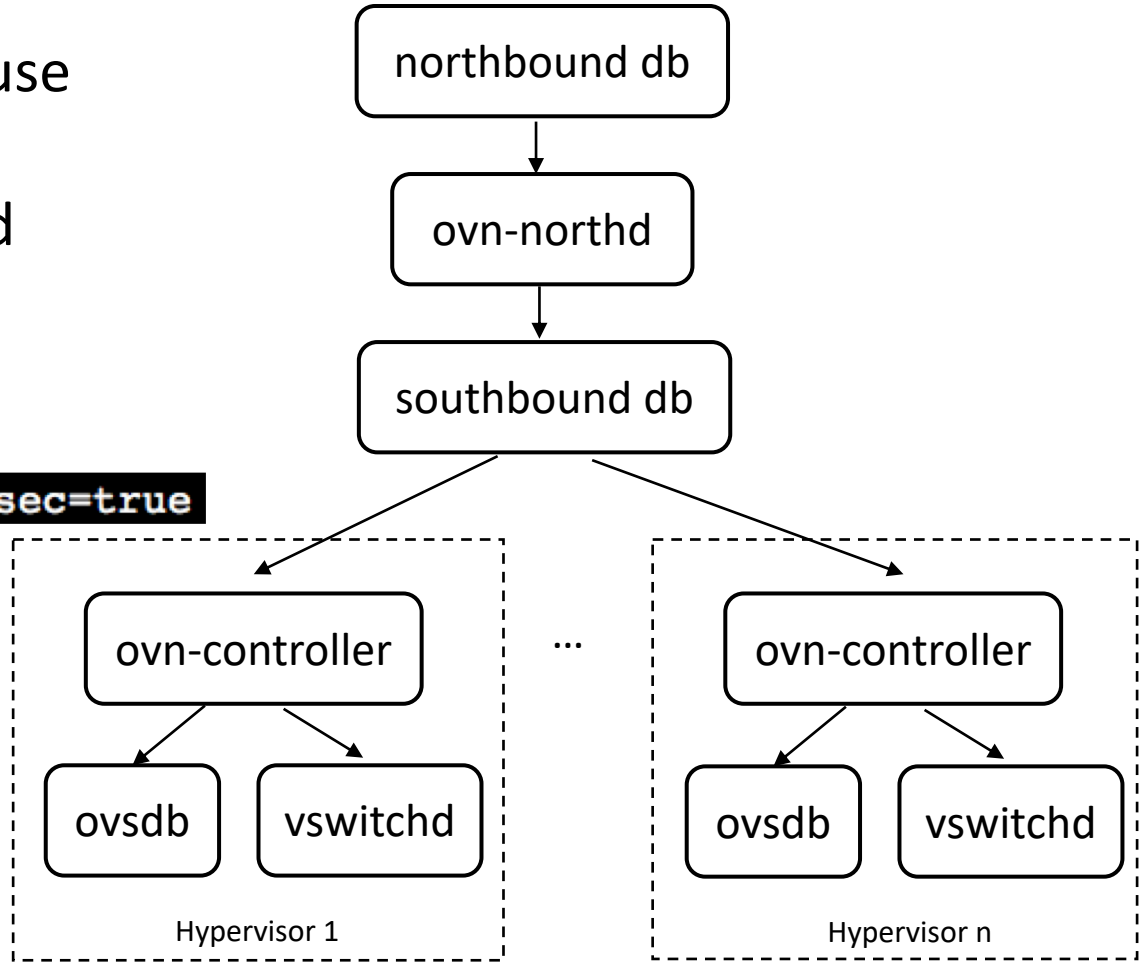

IPsec Evaluation

- Environment: StrongSwan 5.3.5, Linux 4.4.0, Intel Xeon 2 GHz, 10 Gbps NIC
- iperf generates TCP stream (window size: 85KB), which is encrypted in a single core

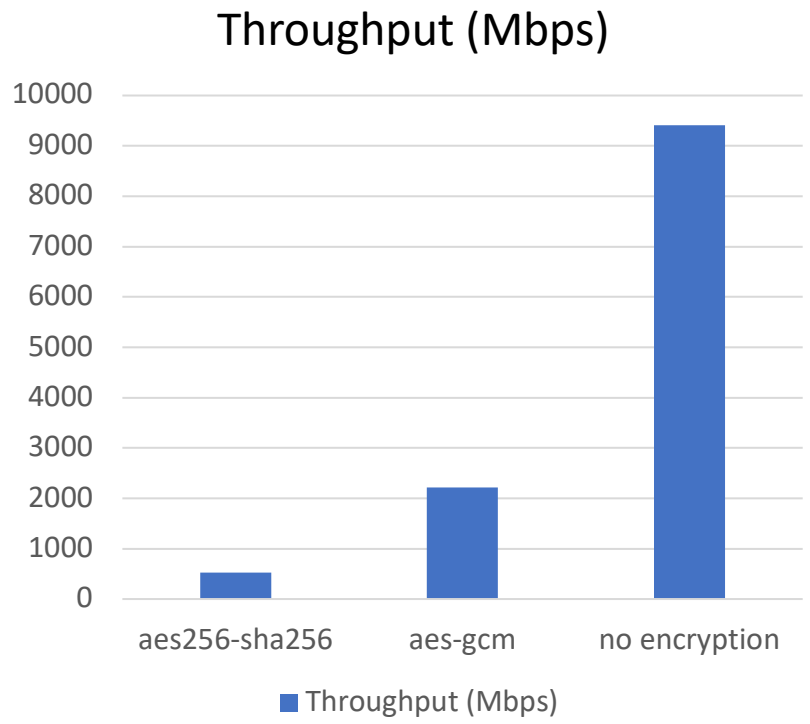

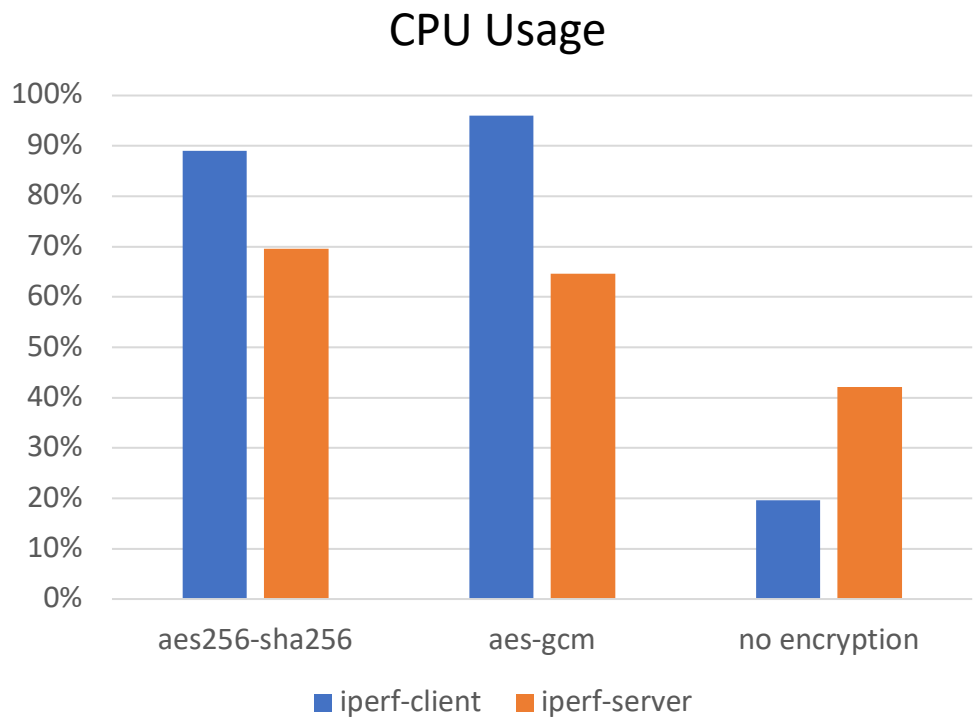

IPsec Evaluation

- Environment: StrongSwan 5.3.5, Linux 4.4.0, Intel Xeon 2 GHz, 10 Gbps NIC
- iperf generates TCP stream (window size: 85KB), which is encrypted in a single core

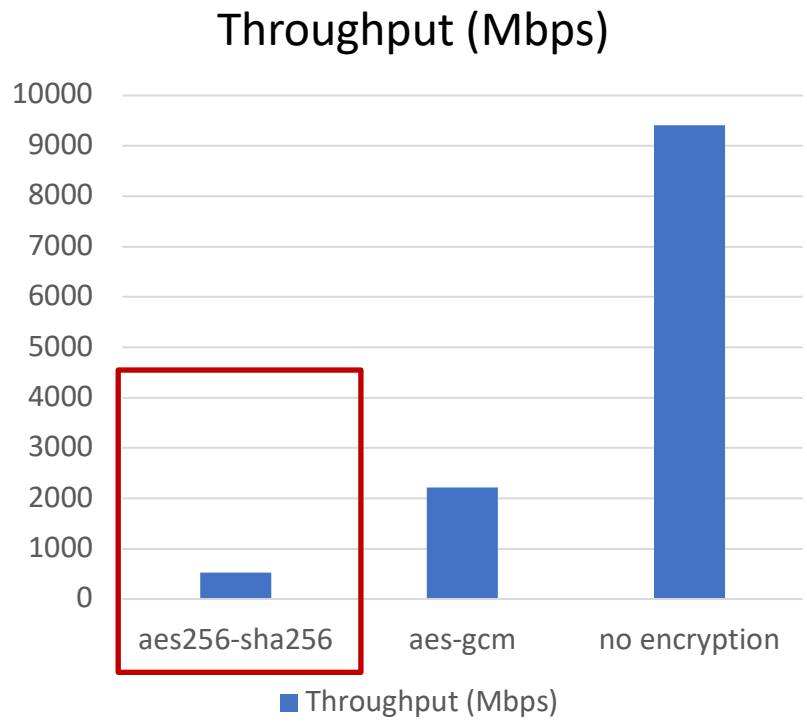

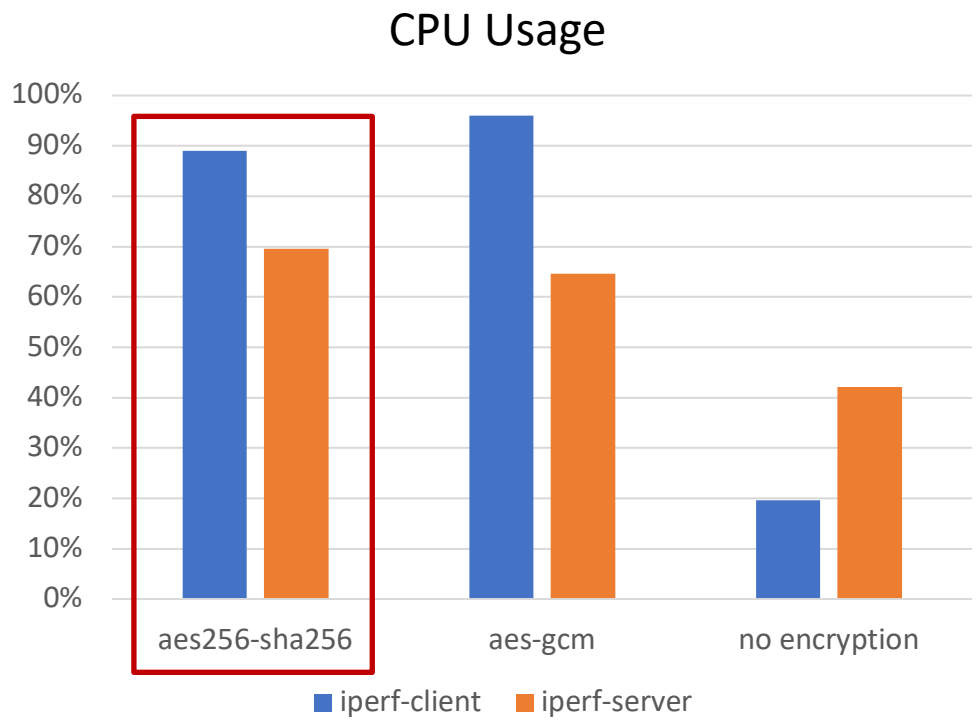

IPsec Evaluation

- Environment: StrongSwan 5.3.5, Linux 4.4.0, Intel Xeon 2 GHz, 10 Gbps NIC
- iperf generates TCP stream (window size: 85KB), which is encrypted in a single core

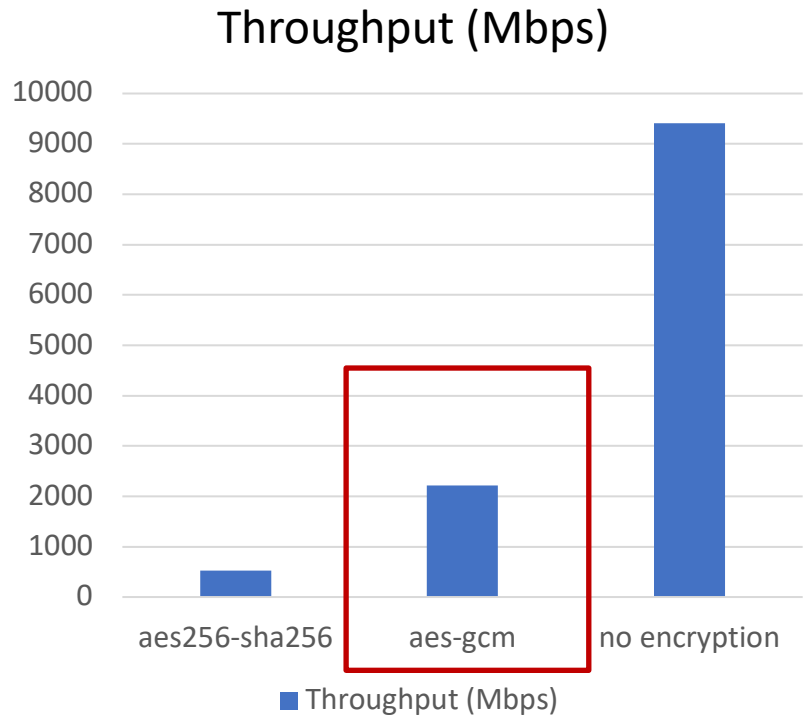

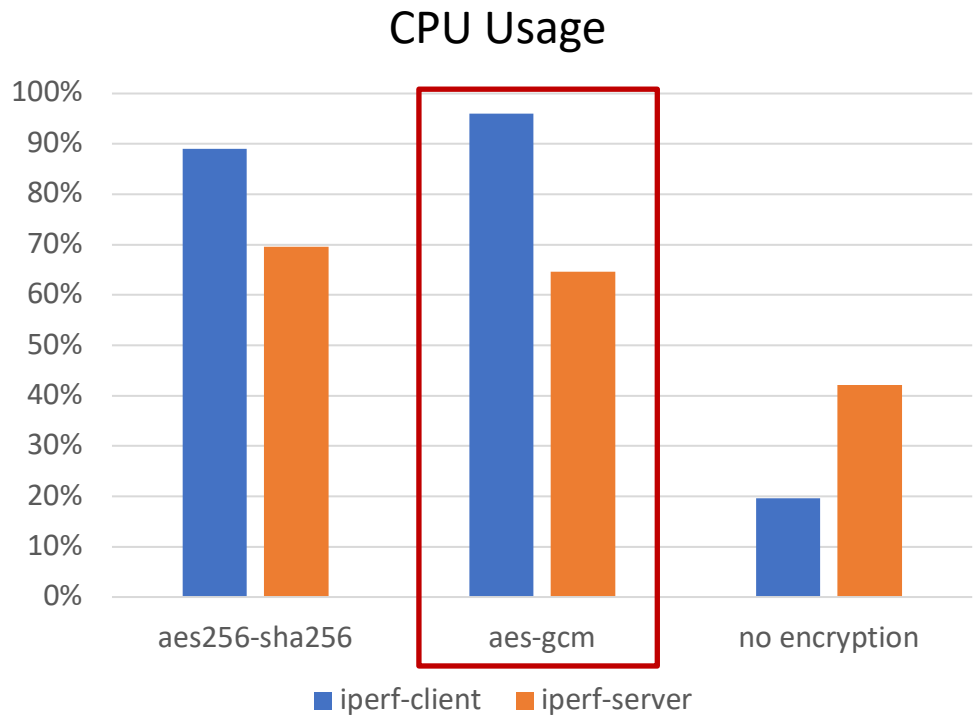

IPsec Evaluation

- Environment: StrongSwan 5.3.5, Linux 4.4.0, Intel Xeon 2 GHz, 10 Gbps NIC
- iperf generates TCP stream (window size: 85KB), which is encrypted in a single core

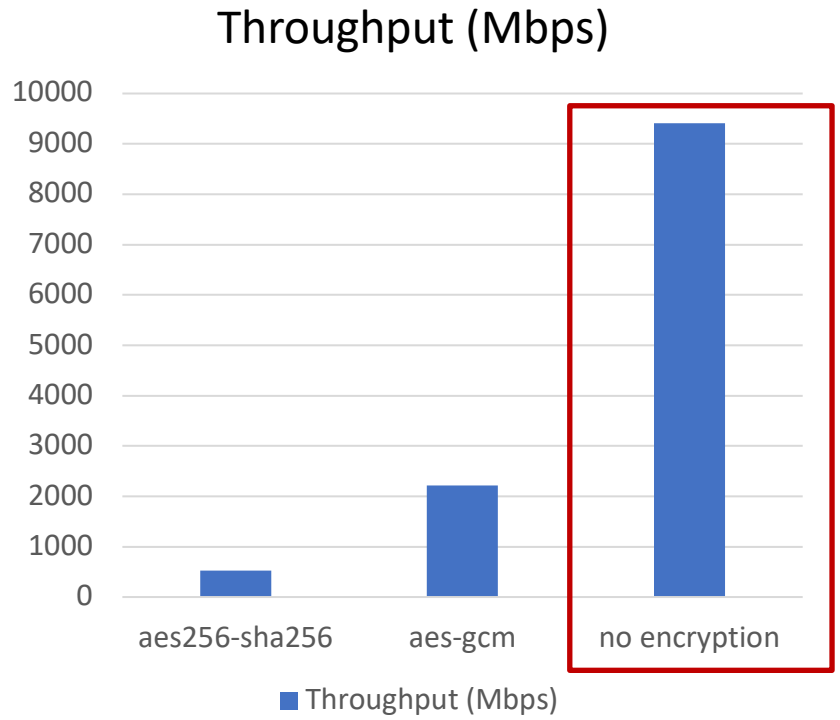

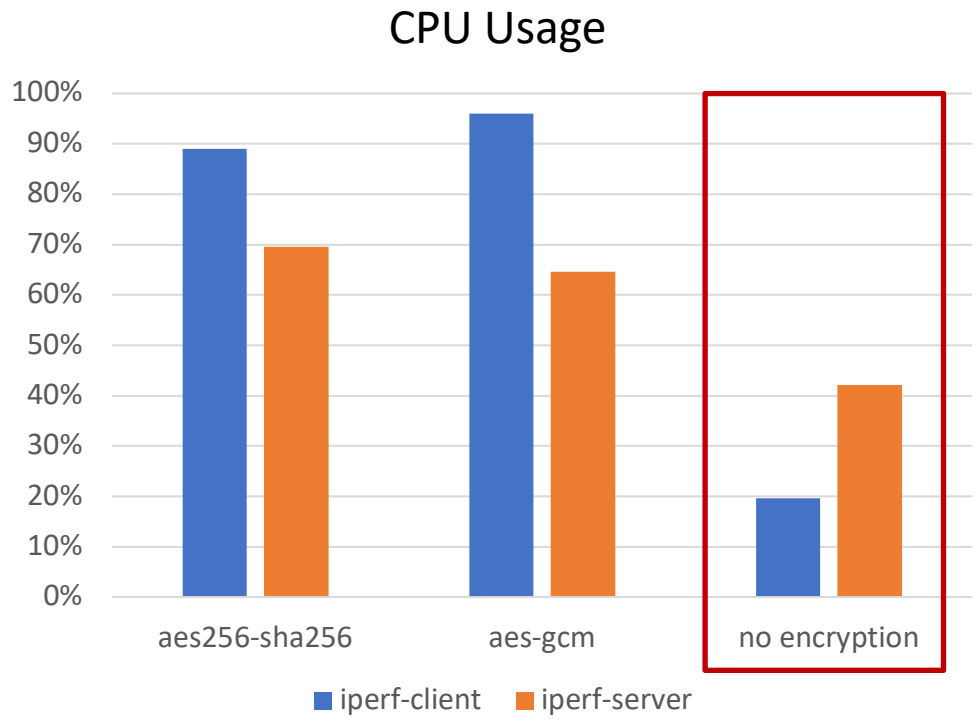

Current Status

- Compatible with StrongSwan and LibreSwan IKE daemon
- Packages for Ubuntu and Fedora
- Tutorials on using OVN IPsec
- Need to use OVS upstream kernel module

Future Directions

More flexible tunnel encryption policies:

- Only encrypting tunnel traffic between certain hypervisors
- Only encrypting tunnel traffic from certain logical network

Q&A