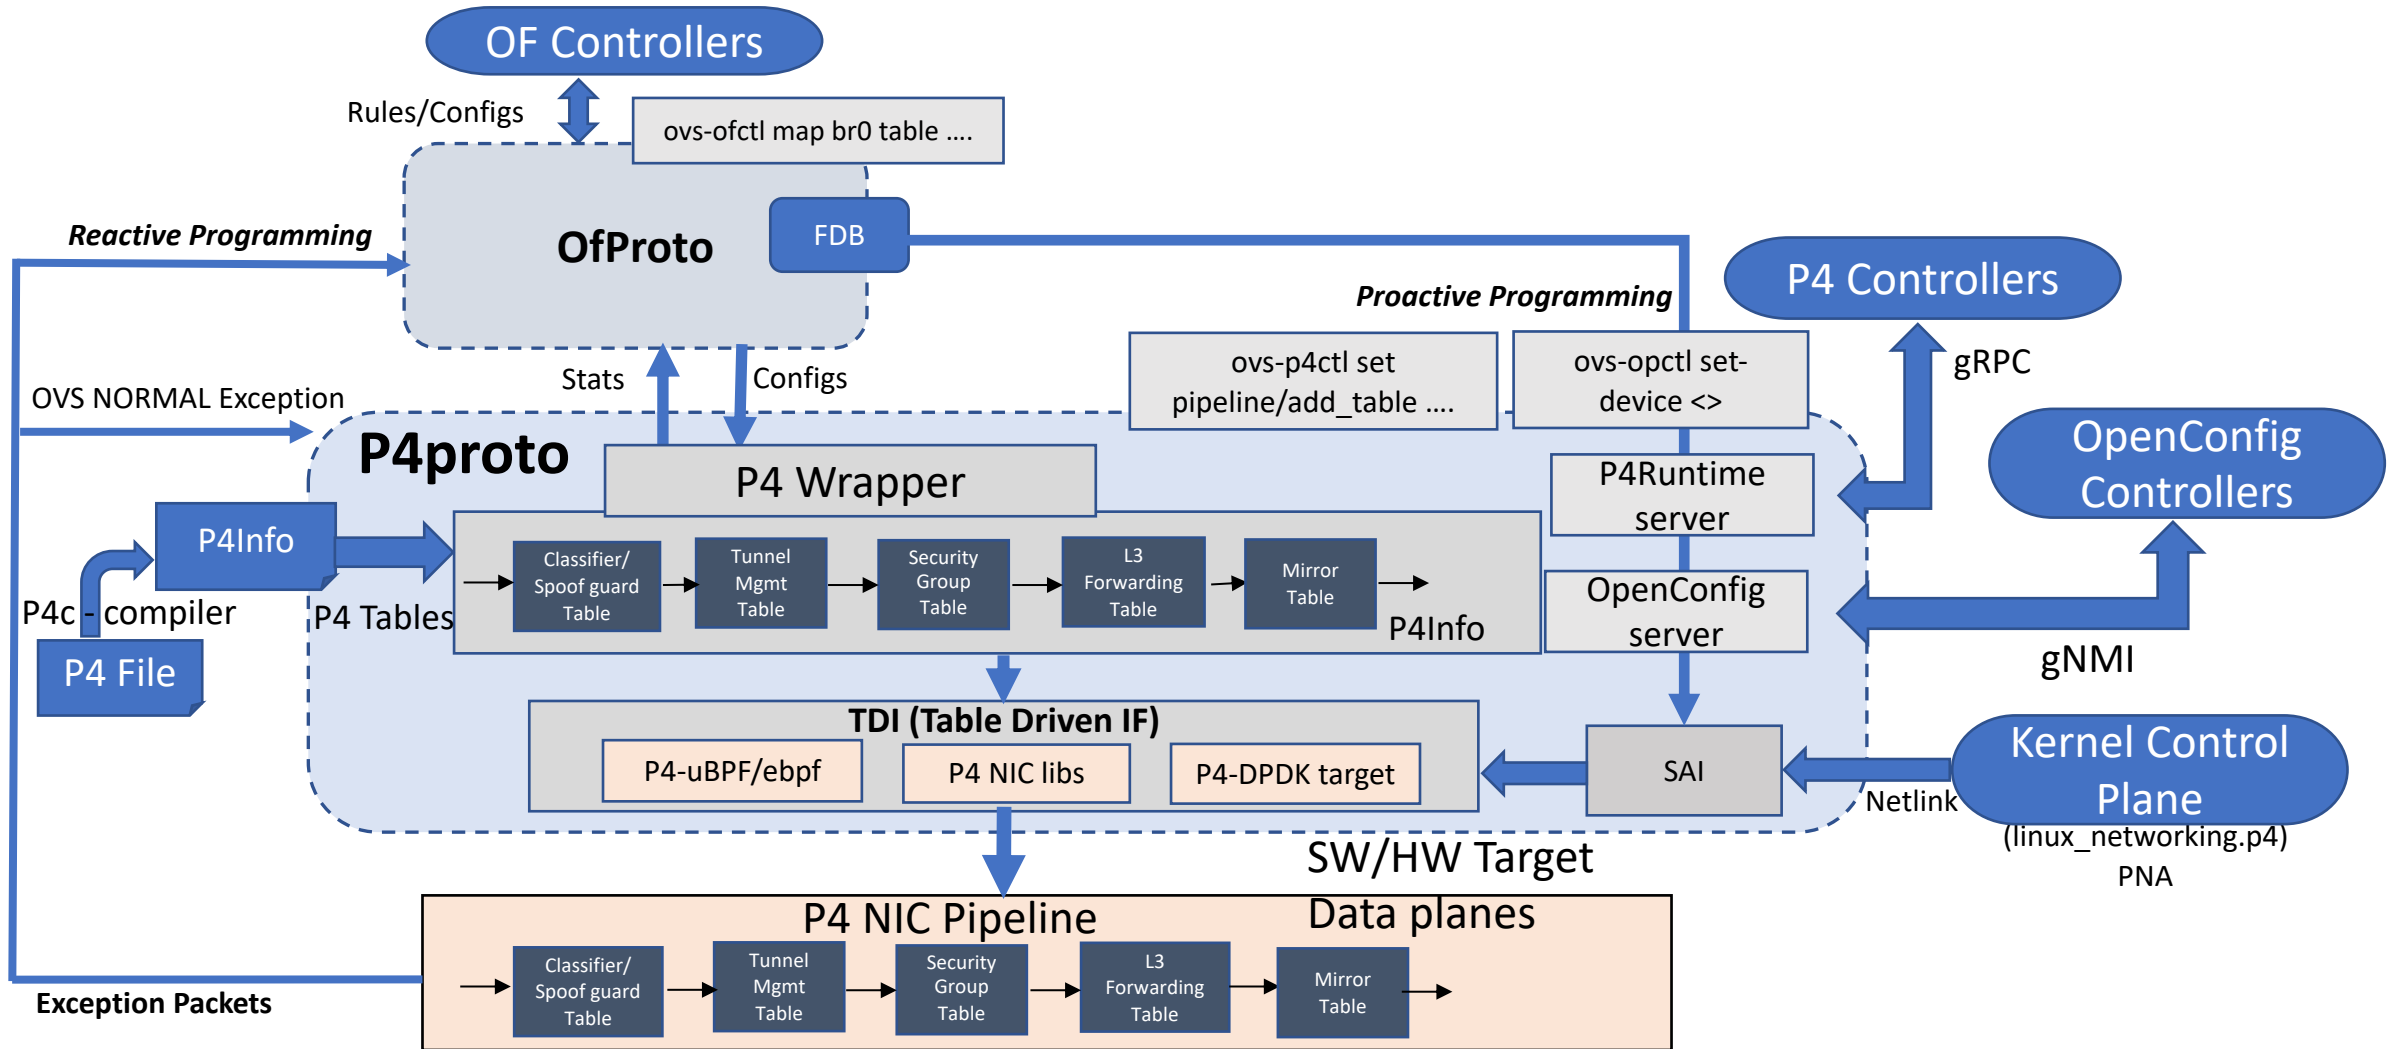




L2 Forwarding and VxLAN with Control Packet Flow using P4-OVS

- Sandeep Nagapattinam
(Cloud Software Engineer, Intel)
- Nupur Uttarwar
(Cloud Software Engineer, Intel)

P4-OVS Architecture



Linux Networking P4 program

Host to Network direction P4 program snippets

```
if (local_metadata.control_packet == 1 && hdr.outer_vlan[0].isValid()) {
    handle_tx_control_vlan_pkts_table.apply();
} else if (local_metadata.control_packet == 1) {
    handle_tx_control_pkts_table.apply();
} else {
    l2_fwd_tx_table.apply();
}
switch (l2_fwd_tx_table.apply().action_run) {
    set_tunnel: {
        ipv4_table.apply();
        nexthop_table.apply();
    }
}
if ((vendormeta_mod_action_ref & (16w1 << VXLAN_ENCAP)) != 0) {
    vxlan_encap_mod_table.apply();
}
switch (neighbor_mod_table.apply().action_run) {
    set_outer_mac: {
        rif_mod_table_start.apply();
        rif_mod_table_mid.apply();
        rif_mod_table_last.apply();
    }
}
```

Network to Host direction P4 program snippets

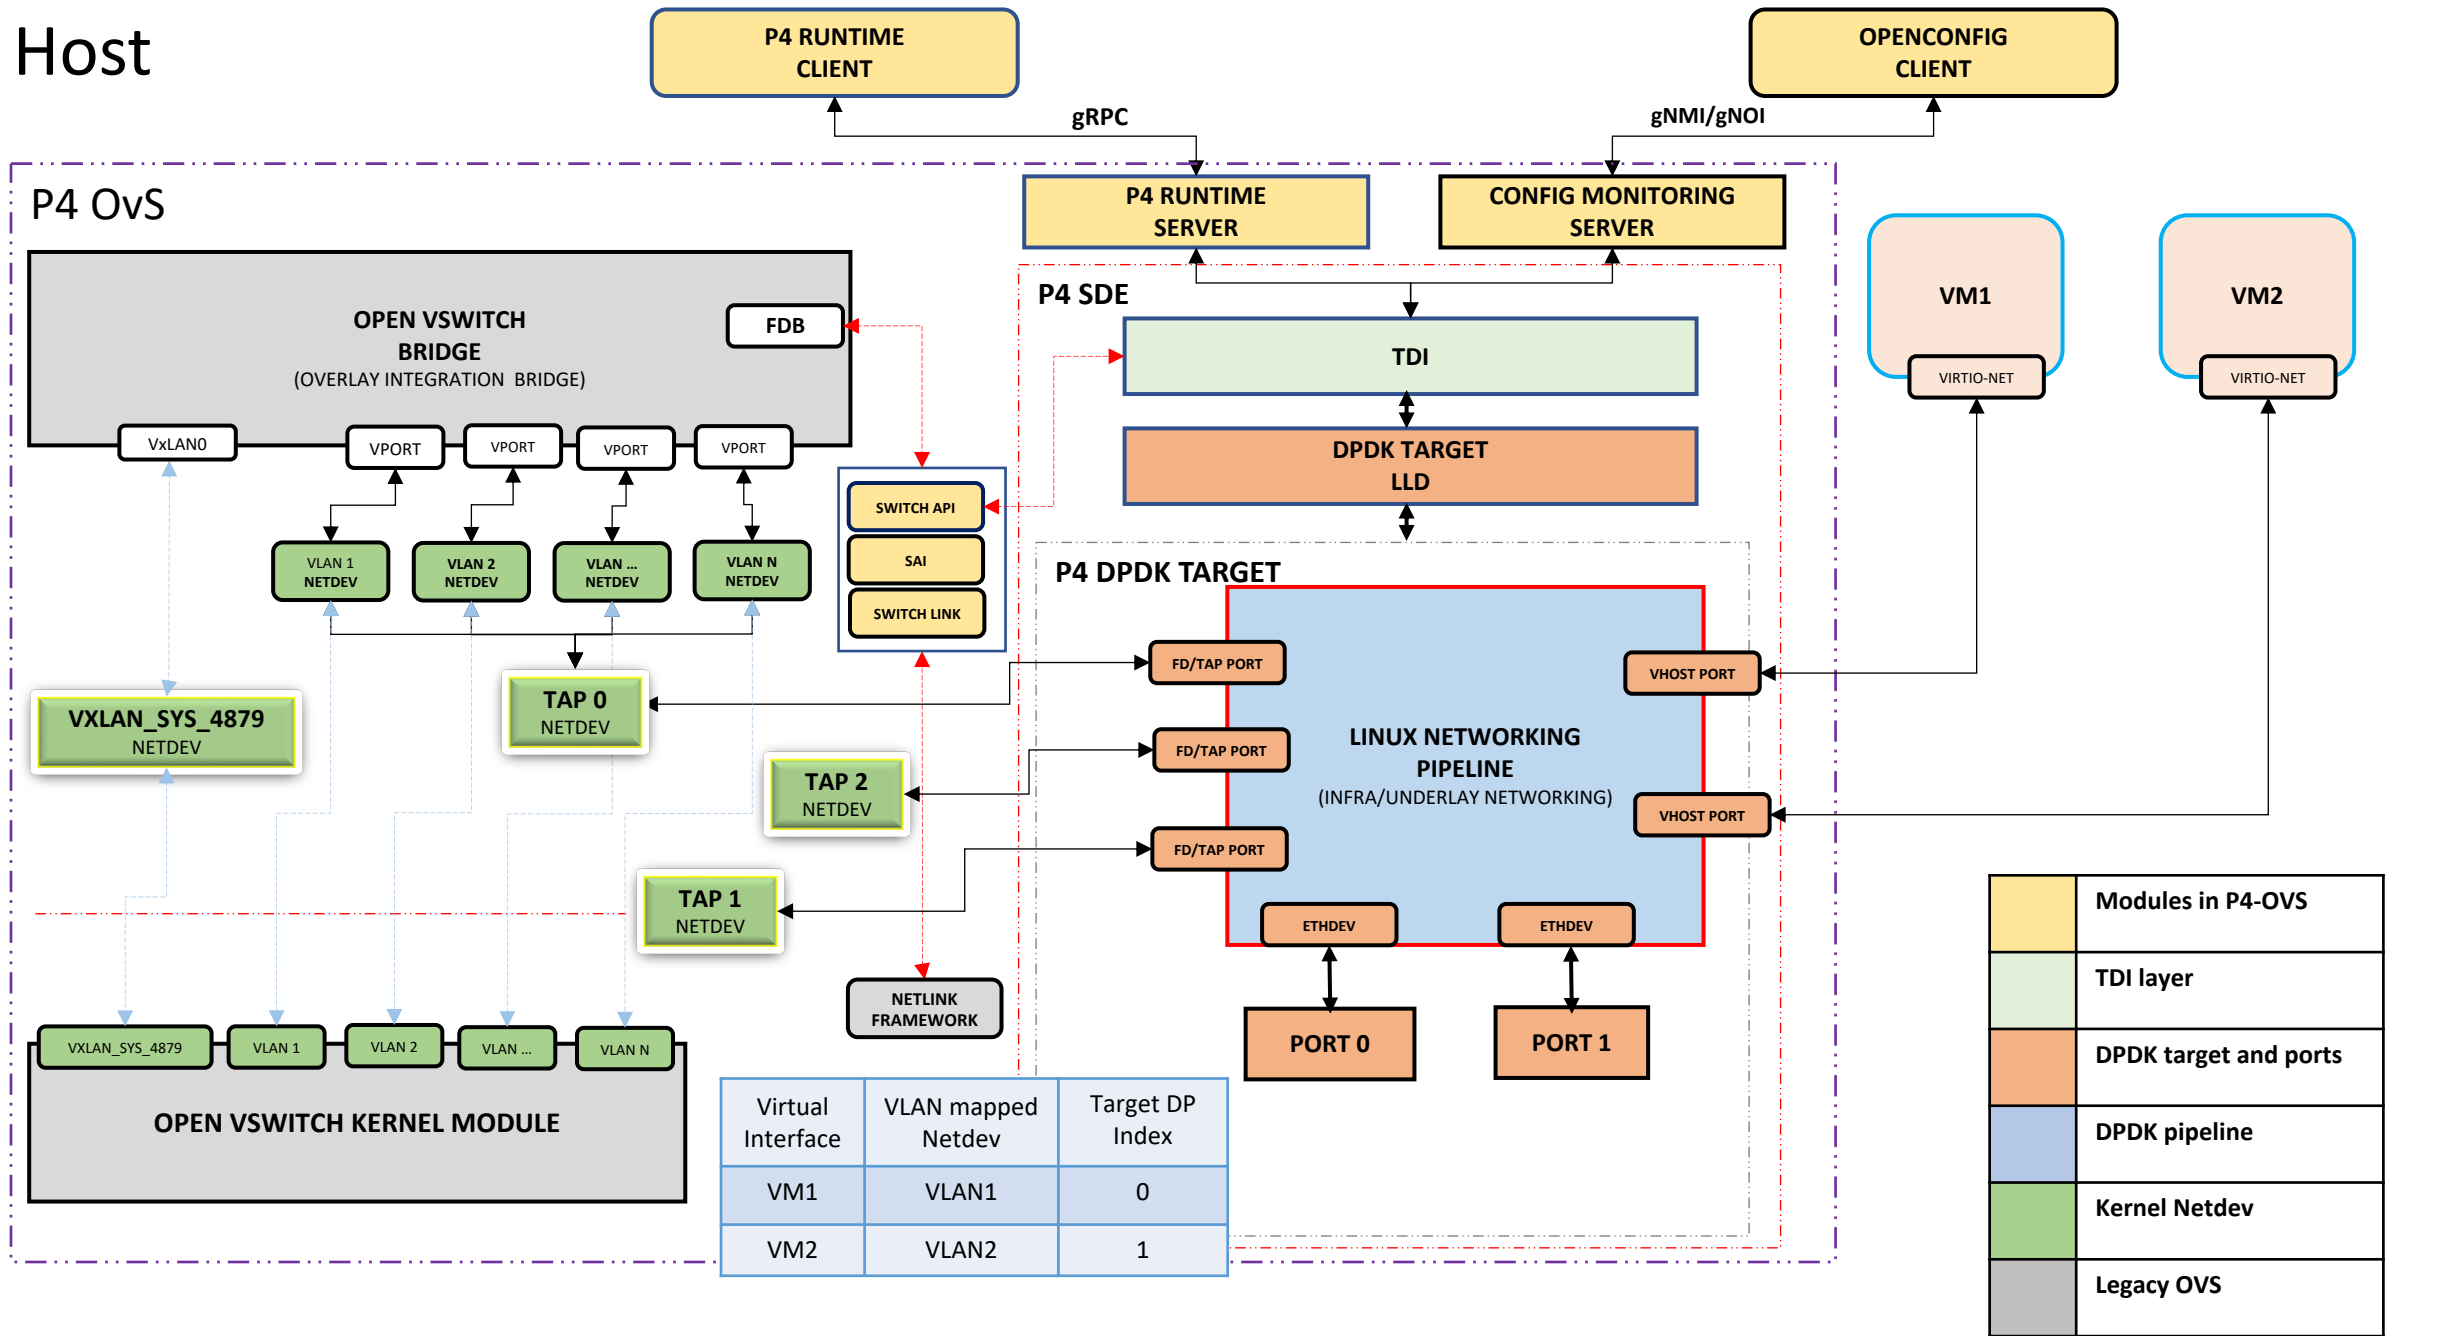
```
if (local_metadata.control_packet == 1) {
    handle_rx_control_pkts_table.apply();
} else if (hdr.outer_ipv4.isValid() && hdr.ethernet.isValid()) {
    ipv4_tunnel_term_table.apply();
    l2_fwd_rx_with_tunnel_table.apply();
} else {
    l2_fwd_rx_table.apply();
}
```

Events and P4 Tables with actions

Event	P4 Table	Action
VxLAN tunnel creation	vxlan_encap_mod_table	vxlan_encap
	ipv4_tunnel_term_table	decap_outer_ipv4
ARP learn by Kernel control plane	neighbor_mod_table	set_outer_mac (DMAC)
	rif_mod_table_start	set_src_mac_start
	rif_mod_table_mid	set_src_mac_mid
	rif_mod_table_last	set_src_mac_last
	l2_fwd_rx_table	l2_fwd
	nexthop_table	set_nexthop
	ipv4_table	set_nexthop_id
FDB learn by OVS control plane	l2_fwd_tx_table	l2_fwd
	l2_fwd_tx_table	set_tunnel
	l2_fwd_rx_with_tunnel_table	l2_fwd

Host

P4 Ovs

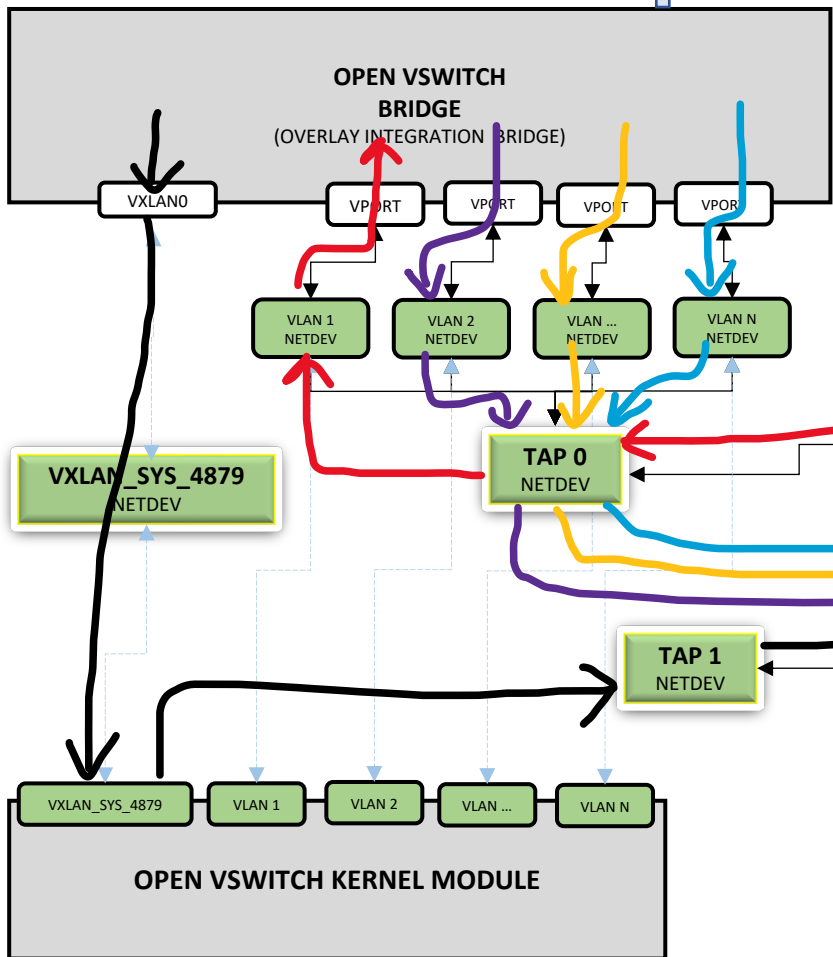


Host

MAC	VLAN mapped Netdev
VM1-MAC	VLAN1

Table	Match Key	Action
I2_fwd_tx_table	VM1-MAC	I2_fwd(vhost-port0)

P4 OvS



P4 SDE

P4 DPDK TARGET

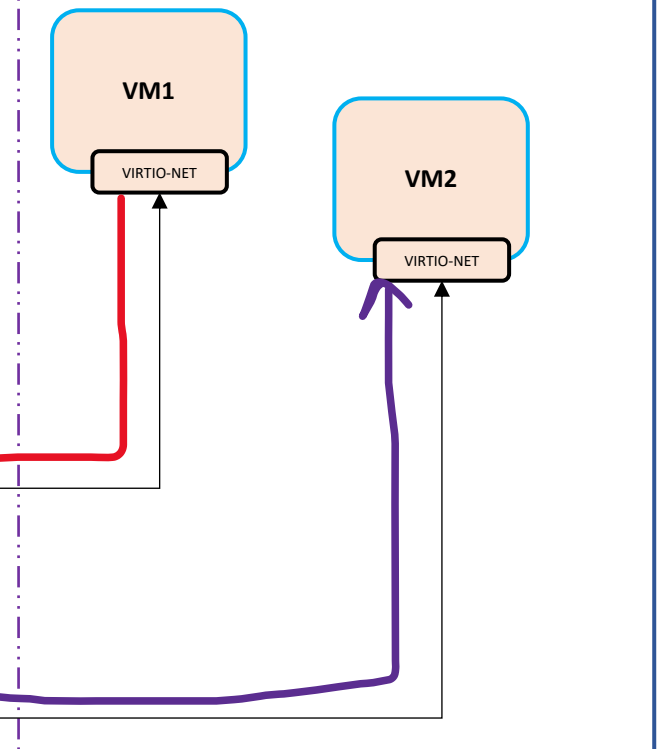
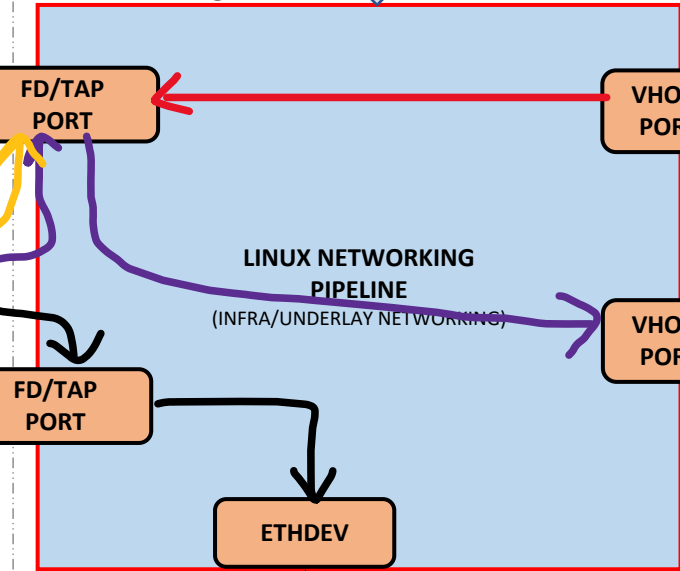


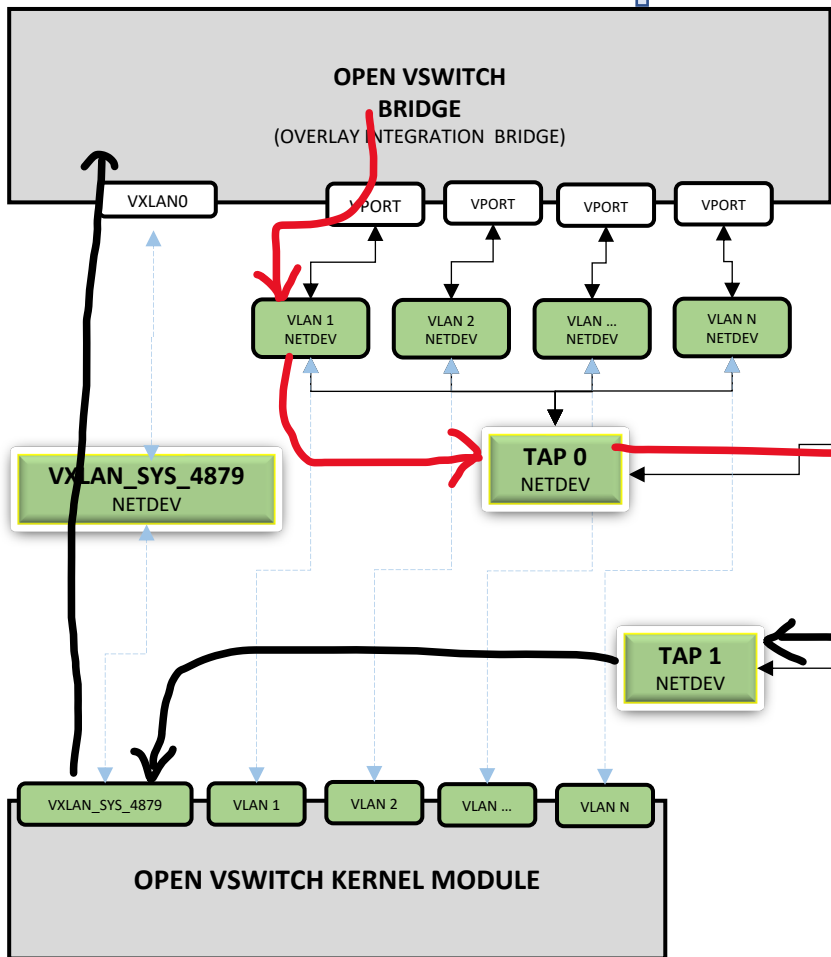
Table	Match Key	Action
I2_fwd_rx_table	TAP1-MAC	I2_fwd(tap1)
ipv4_table	TAP1-IP	set_nexthop_id(id)
nexthop_table	nexthop_id	set_nexthop(id, nbr_id, port)
neighbor_mod_table	nbr_id	set_outer_mac(dmac)

Host

MAC	VLAN mapped Netdev
VM1-MAC	VLAN1
HOST2-VM1-MAC	VxLAN0


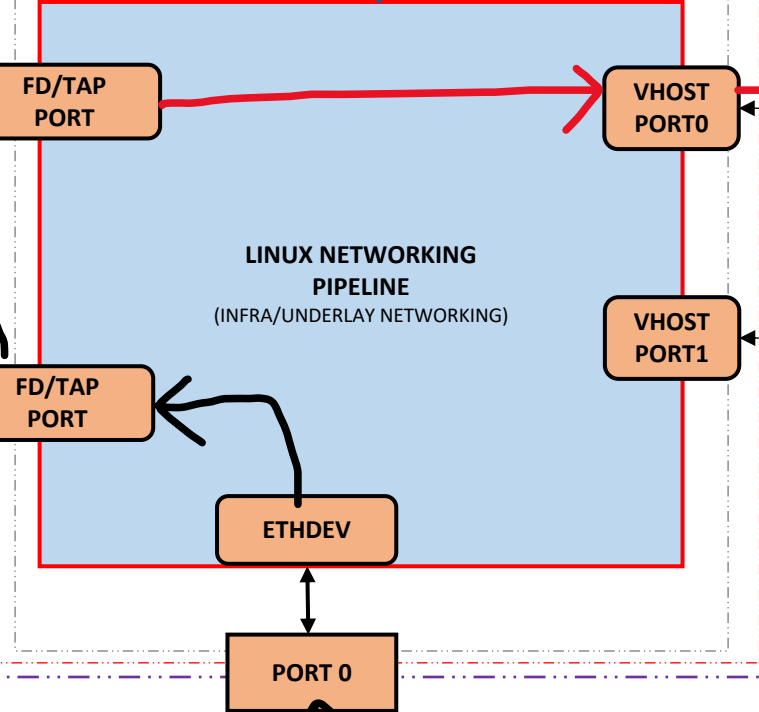
Table	Match Key	Action
l2_fwd_tx_table	VM1-MAC	l2_fwd(vhost-port0)
l2_fwd_tx_table	HOST2-VM1-MAC	set_tunnel(id, dst_addr)

P4 OvS



P4 SDE

P4 DPDK TARGET





GitHub Links:

- <https://github.com/ipdk-io/ovs>
- <https://github.com/p4lang/p4c>
- <https://github.com/opencomputeproject/SAI>

Thank You