

Configuring L2GRE

1, The port ranges in PicOS are as follows:

Port Type	Port Number
L2GRE	5121-6143

PicOS OVS support Layer 2 over Generic Routing Encapsulation (L2GRE); the port number of L2GRE ranges from 5121 to 6143. GRE is an encapsulated mechanism that encapsulates packet IPs; L2GRE is an encapsulated mechanism that encapsulates the entire packet. To resolve the problem that pushes the interface PVID to the untagged packets before encapsulation by the L2GRE header, use the command **ovs-vsctl set interface <interface> type=pica8 options:access-vport=true**. Like this, the untagged packets can be encapsulated by L2GRE header with no VLAN; that is, the PVID of ingress port. And the tagged packets are encapsulated by L2GRE header, the inner VLAN is the VLAN tag of the packets that are received by ingress port. See the example below.

```
admin@PicOS-OVS$ovs-vsctl add-port br0 l2gre1 -- set interface l2gre1 type=pica8_l2gre options:remote_ip=10.10.61.10 options:local_ip=10.10.60.10 options:vlan=1 options:l2gre_key=1234 options:src_mac=C8:0A:A9:9E:49:1A options:dst_mac=C8:0A:A9:9E:14:A5 options:egress_port=qe-1/1/2
```


Description

1. br0: bridge name
2. remote_ip=10.10.61.10: the IP address of the peer L2GRE tunnel interface; this IP address will be the destination IP of the encapsulated L2GRE packets
3. local_ip=10.10.60.10: the IP address of this L2GRE tunnel interface; this IP address will be the source IP of the encapsulated L2GRE packets
4. src_mac=C8:0A:A9:9E:49:1A: the logical MAC address of the L2GRE tunnel interface; this MAC address will be the source MAC of the encapsulated L2GRE packets to next-hop
5. dst_mac=C8:0A:A9:9E:14:A5: the next-hop MAC address; this MAC address will be the destination MAC the encapsulated L2GRE packets to next-hop
6. egress_port=qe-1/1/2: the output port of the encapsulated L2GRE packets
7. l2gre_key=1234:the key value of L2GRE tunnel,different tunnel has different key.
8. vlan=1:the vlan of L2GRE tunnel.this vlan will be pop or not according to the pvid of the egress port.

Examples

push one L2GRE header

topology



Creating a L2GRE tunnel

(1) create a new bridge named br0.

```
admin@PicOS-OVS$ovs-vsctl add-br br0 -- set bridge br0 datapath_type=pica8
```

(2) add ports to br0.


```
admin@PicOS-OVS$ovs-vsctl add-port br0 qe-1/1/1 vlan_mode=trunk tag=1 -- set Interface qe-1/1/1 type=pica8
admin@PicOS-OVS$ovs-vsctl add-port br0 qe-1/1/2 vlan_mode=trunk tag=1 -- set Interface qe-1/1/2 type=pica8
admin@PicOS-OVS$ovs-vsctl add-port br0 l2gre1 -- set interface l2gre1 type=pica8_l2gre options:remote_ip=10.10.61.10 options:local_ip=10.10.60.10 options:vlan=1 options:l2gre_key=1234 options:src_mac=C8:0A:A9:9E:49:1A options:dst_mac=C8:0A:A9:9E:14:A5 options:egress_port=qe-1/1/2
```

User must configure a flow if user wants to send packets to a L2GRE port. And port number is 5121 for l2gre1 tunnel, different L2GRE tunnels must have different l2gre_key.

```
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=1,actions=output:5121
```

Send packets (ARP, L2/L3 packets) to qe-1/1/1, then packets are encapsulated by L2GRE header. When the VLAN tag of Layer 2 GRE tunnel is the same with native VLAN-ID of output port, L2GRE VLAN of the packets are stripped when forwarded by egress port. When the VLAN tag of L2GRE tunnel is different from native VLAN-ID of output port, L2GRE VLAN of the packets are not stripped when forwarded by egress port.

strip L2GRE tunnel



configuration

configure the L2GRE tunnels named l2gre1 on qe-1/1/2 of swa and l2gre2 on qe-1/1/2 of swb.

swa:

```
admin@PicOS-OVS$ovs-vsctl add-port br0 l2gre1 -- set Interface l2gre1 type=pica8_l2gre options:remote_ip=10.10.61.10 options:local_ip=10.10.61.10 options:vlan=2 options:l2gre_key=1234 options:src_mac=C8:0A:A9:04:49:1A options:dst_mac=C8:0A:A9:04:49:1A options:egress_port=qe-1/1/2
```

swb:

```
admin@PicOS-OVS$ovs-vsctl add-port br0 l2gre1 -- set Interface l2gre1 type=pica8_l2gre options:remote_ip=10.10.61.10 options:local_ip=10.10.61.10 options:vlan=2 options:l2gre_key=1234 options:src_mac=C8:0A:A9:04:49:1A options:dst_mac=C8:0A:A9:04:49:1A options:egress_port=qe-1/1/2
```

User must add the two flows below if user wants to push L2GRE header on qe-1/1/2 of swa and strip the Layer 2 GRE header on qe-1/1/2 of swb.

Swa:


```
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=1,actions=output:5121
```

swb:

```
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=5121,actions=output:1
```

qe-1/1/1 of swb will receive the original packets (the contents of packets are the same with packets that qe-1/1/1 of swa received).

configure two L2GRE tunnels on one physical port



Configuration

Configure two L2GRE tunnels on both qe-1/1/2 (l2gre1, l2gre2) and qe-1/1/3 (l2gre1, l2gre2). These two tunnels have different IP and l2gre_key so user must configure some flows.

Sw1:

```
admin@PicOS-OVS$ovs-vsctl add-port br0 l2gre1 -- set Interface l2gre1 type=pica8_l2gre options:remote_ip=10.10.60.10 options:local_ip=10.10.61.10 options:vlan=2 options:l2gre_key=1234 options:src_mac=C8:0A:A9:9E:14:A5 options:dst_mac=C8:0A:A9:04:49:1A options:egress_port=qe-1/1/2
admin@PicOS-OVS$ovs-vsctl add-port br0 l2gre2 -- set Interface l2gre2 type=pica8_l2gre options:remote_ip="10.10.61.61 options:local_ip=10.10.60.60 options:vlan=10 options:l2gre_key=1235 options:src_mac=C8:0A:A9:04:49:1A options:dst_mac=88:88:88:88:88:88 options:egress_port=qe-1/1/2"
```

flows in sw1,

```
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=1,d1_dst=22:66:66:66:66:66,actions=output:5121
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=1,d1_dst=22:66:66:66:66:67,actions=output:5122
```

sw2:

```
admin@PicOS-OVS$ovs-vsctl add-port br0 l2gre1 -- set Interface l2gre1 type=pica8_l2gre options:remote_ip=10.10.61.10 options:local_ip=10.10.60.10 options:vlan=2 options:l2gre_key=1234 options:src_mac=C8:0A:A9:04:49:1A options:dst_mac=C8:0A:A9:9E:14:A5 options:egress_port=qe-1/1/3
admin@PicOS-OVS$ovs-vsctl add-port br0 l2gre2 -- set Interface l2gre2 type=pica8_l2gre options:remote_ip=10.10.60.60 options:local_ip=10.10.61.61 options:vlan=10 options:l2gre_key=1235 options:src_mac=88:88:88:88:88:88 options:dst_mac=C8:0A:A9:04:49:1A options:egress_port=qe-1/1/3
```


flows in sw2

```
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=5121,d1_dst=22:66:66:66:66:66,actions=output:4
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=5122,d1_dst=22:66:66:66:66:67,actions=output:5
```

send packets to qe-1/1/1,different packets will go to different L2GRE tunnels. When they are stripped L2GRE header on qe-1/1/3,they are forwarded to a different port.

Length of l2gre_key

In pica8 switch, the length of l2gre_key can be 16bit, 20bit, 24bit or 32bit; 24 bit is the default value.



configuration

configure the L2GRE tunnel on qe-1/1/2 .


```
admin@PicOS-OVS$ovs-vsctl add-port br0 l2gre1 -- set Interface l2gre1 type=pica8_l2gre options:remote_ip=10.10.61.10 options:local_ip=10.10.60.10 options:vlan=1 options:l2gre_key=1234 options:src_mac=C8:0A:A9:04:49:1A options:dst_mac=C8:0A:A9:9E:14:A5 options:egress_port=qe-1/1/2
```

Add a flow to switch

```
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=1,actions=output:5121
```

This key of L2GRE tunnel is 1234 here in decimal, 4d2 in hex. The default value of Layer 2 GRE key is 20, so the value of GRE key of packets is 0x004d2000. When user sets the l2gre_key value to 16 using the command **ovs-vsctl set-l2gre-key-length 16**, the value of the GRE key of packet is 0x04d20000. When user sets the l2gre_key value to 20 using the command **ovs-vsctl set-l2gre-key-length 24**, the value of GRE key packet is 0x0004d200. When user sets the l2gre_key value to 32 using the command **ovs-vsctl set-l2gre-key-length 32**, the value of GRE key packet is 0x000004d2.

Collaboration between nvgre and VXLAN



Configuration

configure the L2GRE tunnel and VXLAN tunnel on qe-1/1/2 and qe-1/1/3.

```

admin@PicOS-OVS$ovs-vsctl add-port br0 l2gre1 -- set Interface l2gre1 type=pica8_l2gre options:remote_ip=10.10.61.10 options:local_ip=10.10.60.10 options:vlan=1 options:l2gre_key=1234 options:src_mac=C8:0A:A9:04:49:1A options:dst_mac=C8:0A:A9:9E:14:A5 options:egress_port=qe-1/1/2
admin@PicOS-OVS$
admin@PicOS-OVS$ovs-vsctl add-port br0 l2gre1 -- set Interface l2gre1 type=pica8_l2gre options:remote_ip=10.10.60.10 options:local_ip=10.10.61.10 options:vlan=1 options:l2gre_key=1234 options:src_mac=C8:0A:A9:9E:14:A5 options:dst_mac=C8:0A:A9:04:49:1A options:egress_port=qe-1/1/3
admin@PicOS-OVS$ovs-vsctl add-port br0 vxlan1 -- set interface vxlan1 type=pica8_vxlan options:remote_ip=10.10.10.2 options:local_ip=10.10.10.1 options:vlan=1 options:vnid=1122867 options:udp_dst_port=4789 options:src_mac=66:66:66:77:77:77 options:dst_mac=88:88:88:77:77:77 options:egress_port=qe-1/1/2
admin@PicOS-OVS$
admin@PicOS-OVS$ovs-vsctl add-port br0 vxlan1 -- set interface vxlan1 type=pica8_vxlan options:remote_ip=10.10.10.1 options:local_ip=10.10.10.2 options:vlan=1 options:vnid=1122867 options:udp_dst_port=4789 options:src_mac=88:88:88:77:77:77 options:dst_mac=66:66:66:77:77:77 options:egress_port=qe-1/1/3
  
```

Flows in Switches

sw1:

```

admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=1,dl_dst=22:22:22:22:22:23,actions=output:4097
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=1,dl_dst=22:22:22:22:22:22,actions=output:5121
  
```

sw2:

```

admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=5122,dl_dst=22:22:22:22:22:22,actions=output:4
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=4098,dl_dst=22:22:22:22:22:23,actions=output:5
  
```

qe-1/1/4 should receive the de-capsulated packets with src_mac 22:22:22:22:22:22, qe-1/1/5 should receive the de-capsulated packets with src_mac 22:22:22:22:22:23. That is to say VXLAN and L2GRE do not affect each other.