


Configuring NAT flow

The Network Address Translation (NAT) process maps IP addresses from one address domain (or realm) to another to provide transparent routing to end hosts. Typically, NAT allows organizations to map public external addresses to private or unregistered addresses. Platforms with ASIC Trident2, Trident2Plus and Tomahawk support this function in OVS mode only. A flow with NAT actions (changing IP address or L4 port) can be hardware switched. Flows can be associated with the following actions: `mod_nw_dst`, `mod_nw_src`, `mod_tp_dst` and `mod_tp_src`.

Listed below is the minimal information needed to process the packet on hardware only (direct flow):

- 1) `dl_dst`(match field or action), `dl_vlan`(match field or action), `mod_nw_src`, `mod_tp_src`
- 2) `dl_dst`(match field or action), `dl_vlan`(match field or action), `mod_nw_dst`, `mod_tp_dst`
- 3) `dl_dst`(match field or action), `dl_vlan`(match field or action), `tp_src`(with or without in match field), `mod_nw_src`
- 4) `dl_dst`(match field or action), `dl_vlan`(match field or action), `nw_src`(match field), `mod_tp_src`
- 5) `dl_dst`(match field or action), `dl_vlan`(match field or action), `tp_dst`(with or without in match field), `mod_nw_dst`
- 6) `dl_dst`(match field or action), `dl_vlan`(match field or action), `nw_dst`(match field), `mod_tp_dst`



Example 1: SNAT

Step 1: Create a new bridge named br0.

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

Step 2: Add ports to br0.

```
admin@PicOS-OVS$ovs-vsctl add-port br0 te-1/1/1 vlan_mode=trunk tag=1-- set interface te-1/1/1 type=pica8
admin@PicOS-OVS$ovs-vsctl add-port br0 te-1/1/2 vlan_mode=trunk tag=1-- set interface te-1/1/2 type=pica8
```

Step 3: If user is inside network and wants to visit outside network, source IP needs to be modified:

```
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=1,tcp,dl_vlan=1999,dl_dst=22:22:22:22:22:22,actions=set_field:192.168.5.5->nw_src,set_field:888->tp_src,output:2
```

Step 4: Check flow tables.

```
admin@PicOS-OVS$ovs-ofctl dump-flows br0
OFPST_FLOW reply (OF1.4) (xid=0x2):
 cookie=0x0, duration=264.202s, table=0, n_packets=n/a, n_bytes=0, tcp,in_port=1,dl_vlan=1999dl_dst=22:22:22:22:22:22 actions=set_field:192.168.5.5->ip_src,set_field:888->tcp_src,output:2
admin@PicOS-OVS$
admin@PicOS-OVS$ovs-appctl pica/dump-flows
#7417 normal permanent recirc_id=0,tcp,in_port=1,dl_vlan=1999,dl_dst=22:22:22:22:22:22,actions:set(ipv4(src=192.168.5.5)),set(tcp(src=888)),2
#7416 normal permanent priority=0,recirc_id=0, actions:drop
Total 2 flows in HW.
admin@PicOS-OVS$
```

Example 2: DNAT

Step 1: Create a new bridge named br0

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

Step 2: Add ports to br0

```
admin@PicOS-OVS$ovs-vsctl add-port br0 te-1/1/1 vlan_mode=trunk tag=1-- set interface te-1/1/1 type=pica8
admin@PicOS-OVS$ovs-vsctl add-port br0 te-1/1/2 vlan_mode=trunk tag=1-- set interface te-1/1/2 type=pica8
```

Step 3: If user is outside network and wants to visit inside network, destination IP needs to be modified:

```
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=1,tcp,dl_vlan=1999,dl_dst=22:22:22:22:22:22,actions=set_field:
192.168.5.5->\nw_dst,set_field:800->\>tp_dst,output:2
```

Step 4: Check flow tables

```
admin@PicOS-OVS$ovs-ofctl dump-flows br0
OFPST_FLOW reply (OF1.4) (xid=0x2):
  cookie=0x0, duration=264.202s, table=0, n_packets=n/a, n_bytes=0, tcp,in_port=1,dl_vlan=1999,dl_dst=22:22:22:
22:22:22 actions=set_field:192.168.5.5->ip_dst,set_field:888->tcp_dst,output:2
admin@PicOS-OVS$
admin@PicOS-OVS$ovs-appctl pica/dump-flows
#7417 normal permanent recirc_id=0,tcp,in_port=1,dl_vlan=1999,dl_dst=22:22:22:22:22:22,actions:set(ipv4(dst=192.
168.5.5)),set(tcp(dst=888)),2
#7416 normal permanent priority=0,recirc_id=0, actions:drop
Total 2 flows in HW.
admin@PicOS-OVS$
```

Example 3: Packet-driven-flow

If match field or actions cannot satisfy condition of direct flow, this flow will be packet-driven-flow, and it cannot be added to hardware table directly.

Establish br0 and add ports in br0 like above configuration. And add flow as follows:

```
admin@PicOS-OVS$ovs-ofctl add-flow br0 in_port=1,dl_type=0x0800,tcp,actions=set_field:192.168.5.5->\nw_src,
set_field:0x456->\>tp_src,output:2
admin@PicOS-OVS$
admin@PicOS-OVS$ovs-ofctl dump-flows br0
OFPST_FLOW reply (OF1.4) (xid=0x2):
  cookie=0x0, duration=3.442s, table=0, n_packets=n/a, n_bytes=0, tcp,in_port=1 actions=set_field:192.168.5.5-
>ip_src,set_field:1110->tcp_src,output:2
admin@PicOS-OVS$ovs-appctl pica/dump-flows
#7460 normal permanent recirc_id=0,tcp,in_port=1, actions:To_CPU(for_packet_driven)
#7458 normal permanent priority=0,recirc_id=0, actions:drop
Total 2 flows in HW.
admin@PicOS-OVS$
```

Sending increasing dst_mac packets to te-1/1/1, mac address from 22:22:22:22:22:22 to 22:22:22:22:22:2b, then check tables:

```

admin@PicOS-OVS$ovs-ofctl dump-flows br0
OFPST_FLOW reply (OF1.4) (xid=0x2):
  cookie=0x0, duration=135.680s, table=0, n_packets=n/a, n_bytes=1124112312, tcp,in_port=1 actions=set_field:
  192.168.5.5->ip_src,set_field:1110->tcp_src,output:2
admin@PicOS-OVS$
admin@PicOS-OVS$ovs-appctl pica/dump-flows
#7479 normal priority=1048560,recirc_id=0,tcp,in_port=1,vlan_tci=0x0000/0x1fff,dl_dst=22:22:22:22:22:22,
nw_src=2.0.0.2,nw_frag=no,tp_src=18,tp_dst=35, actions:set(ipv4(src=192.168.5.5)),set(tcp(src=1110)),2
#7478 normal priority=1048560,recirc_id=0,tcp,in_port=1,vlan_tci=0x0000/0x1fff,dl_dst=22:22:22:22:22:2a,
nw_src=2.0.0.2,nw_frag=no,tp_src=18,tp_dst=35, actions:set(ipv4(src=192.168.5.5)),set(tcp(src=1110)),2
#7477 normal priority=1048560,recirc_id=0,tcp,in_port=1,vlan_tci=0x0000/0x1fff,dl_dst=22:22:22:22:22:29,
nw_src=2.0.0.2,nw_frag=no,tp_src=18,tp_dst=35, actions:set(ipv4(src=192.168.5.5)),set(tcp(src=1110)),2
#7475 normal priority=1048560,recirc_id=0,tcp,in_port=1,vlan_tci=0x0000/0x1fff,dl_dst=22:22:22:22:22:27,
nw_src=2.0.0.2,nw_frag=no,tp_src=18,tp_dst=35, actions:set(ipv4(src=192.168.5.5)),set(tcp(src=1110)),2
#7471 normal priority=1048560,recirc_id=0,tcp,in_port=1,nw_src=192.168.5.5,nw_frag=no,tp_src=1110, actions:2
#7476 normal priority=1048560,recirc_id=0,tcp,in_port=1,vlan_tci=0x0000/0x1fff,dl_dst=22:22:22:22:22:28,
nw_src=2.0.0.2,nw_frag=no,tp_src=18,tp_dst=35, actions:set(ipv4(src=192.168.5.5)),set(tcp(src=1110)),2
#7472 normal priority=1048560,recirc_id=0,tcp,in_port=1,vlan_tci=0x0000/0x1fff,dl_dst=00:00:00:00:00:00,
nw_src=2.0.0.2,nw_frag=no,tp_src=18,tp_dst=35, actions:set(ipv4(src=192.168.5.5)),set(tcp(src=1110)),2
#7484 normal permanent recirc_id=0,tcp,in_port=1, actions:To_CPU(for_packet_driven)
#7483 normal permanent priority=0,recirc_id=0, actions:drop
#7480 normal priority=1048560,recirc_id=0,tcp,in_port=1,vlan_tci=0x0000/0x1fff,dl_dst=22:22:22:22:22:23,
nw_src=2.0.0.2,nw_frag=no,tp_src=18,tp_dst=35, actions:set(ipv4(src=192.168.5.5)),set(tcp(src=1110)),2
#7482 normal priority=1048560,recirc_id=0,tcp,in_port=1,vlan_tci=0x0000/0x1fff,dl_dst=22:22:22:22:22:26,
nw_src=2.0.0.2,nw_frag=no,tp_src=18,tp_dst=35, actions:set(ipv4(src=192.168.5.5)),set(tcp(src=1110)),2
#7481 normal priority=1048560,recirc_id=0,tcp,in_port=1,vlan_tci=0x0000/0x1fff,dl_dst=22:22:22:22:22:24,
nw_src=2.0.0.2,nw_frag=no,tp_src=18,tp_dst=35, actions:set(ipv4(src=192.168.5.5)),set(tcp(src=1110)),2
#7474 normal priority=1048560,recirc_id=0,tcp,in_port=1,vlan_tci=0x0000/0x1fff,dl_dst=22:22:22:22:22:2b,
nw_src=2.0.0.2,nw_frag=no,tp_src=18,tp_dst=35, actions:set(ipv4(src=192.168.5.5)),set(tcp(src=1110)),2
#7473 normal priority=1048560,recirc_id=0,tcp,in_port=1,vlan_tci=0x0000/0x1fff,dl_dst=22:22:22:22:22:25,
nw_src=2.0.0.2,nw_frag=no,tp_src=18,tp_dst=35, actions:set(ipv4(src=192.168.5.5)),set(tcp(src=1110)),2
Total 14 flows in HW.
admin@PicOS-OVS$

```

1. And te-1/1/2 receive packets with the src_ip 192.168.5.5 and src_port 1110.

Due to ASIC limitation, a flow can not modify I4_src_port without modifying SIP or modify I4_dst_port without modifying DIP.

If only modifying SIP(DIP) or SIP+L4_SRC_PORT(DIP+L4_DST_PORT), up to 2k flow can be configured. If modifying both SIP[L4_SRC_PORT] and DIP[L4_DST_PORT], the flow supported is 1k.

2. If set_dl_src is included in actions, the packets will be stamped with set_dl_src (as before). If set_dl_src is not included in actions, the packets will be stamped with the original dl_src. That is to say, keep the original source mac address.

```
Eg1:
ovs-ofctl add-flow br0 in_port=1,ip,tcp,dl_vlan=2,dl_src=00:11:22:33:44:55,dl_dst=00:01:02:03:04:05,
actions=set_field:0x456->tp_src,set_field:192.168.5.5->nw_src,output:2
send packets
result:

MAC: ----- MAC Header -----
MAC: Destination Address : 00 01 02 03 04 05
MAC: Source Address : 00 11 22 33 44 55

Eg2:
ovs-ofctl add-flow br0 in_port=1,ip,tcp,dl_vlan=2,dl_dst=00:01:02:03:04:05,actions=set_field:0x456->tp_src,
set_field:192.168.5.5->nw_src,output:2
send packets
result:

MAC: ----- MAC Header -----
MAC: Destination Address : 00 01 02 03 04 05
MAC: Source Address : 00 11 22 33 44 55

Eg3:
ovs-ofctl add-flow br0 in_port=1,ip,tcp,dl_vlan=2,dl_src=00:11:22:33:44:55,dl_dst=00:01:02:03:04:05,
actions=set_field:0x456->tp_src,set_field:192.168.5.5->nw_src,set_field:22:22:22:22:22:22->dl_src,output:2
send packets
result:

MAC: ----- MAC Header -----
MAC: Destination Address : 00 01 02 03 04 05
MAC: Source Address : 22 22 22 22 22 22

Eg4:
ovs-ofctl add-flow br0 in_port=1,ip,tcp,dl_vlan=2,dl_dst=00:01:02:03:04:05,actions=set_field:0x456->tp_src,
set_field:192.168.5.5->nw_src,set_field:22:22:22:22:22:22->dl_src,output:2
send packets
result:

MAC: ----- MAC Header -----
MAC: Destination Address : 00 01 02 03 04 05
MAC: Source Address : 22 22 22 22 22 22
```