

# EVPN Symmetric Routing Configuration Example

## Note



EVPN feature is currently supported on X86 platforms only.

The following steps represent the fundamental configuration to use EVPN as the control plane for VXLAN. These steps are in addition to configuring VXLAN interfaces, attaching them to a bridge, and mapping VLANs to VNIs.

1. Configure Physical interfaces and assign VLANs to interfaces.
2. Configure L3 interfaces and assign IP addresses to interfaces.
3. Configure VXLAN VNIs and enable VXLAN VNI mapping to VLAN IDs.
4. Enable EVPN route exchange (that is, address-family layer 2 VPN/EVPN) between BGP peers.
5. Enable EVPN on the system to advertise VNIs and host reachability information (MAC addresses learned on associated VLANs) to BGP peers.

Additional configuration is necessary such as provision of inter-subnet routing. The configuration depends on the deployment scenario. You can also configure various other BGP parameters depending on your network requirements.

## Enable EVPN Between BGP Peers

The very basic steps needed to enable BGP EVPN between a BGP neighbor is as under.

```
admin@router1# set protocols bgp local-as 65101
admin@router1# set protocols bgp router-id 10.10.10.1
admin@router1# set protocols bgp neighbor 100.1.1.134 remote-as 650101
admin@router1# commit
Commit OK.
Save Done!
```

The configuration below adds the evpn address family to the BGP neighbor address-family so that BGP peers activate exchanging EVPN routes with each other. After this configuration, the BGP still does not know about the local VNI's.

## Advertise All VNIs Through BGP

To allow BGP know about all VNIs or hosts associated with those local VNIs, enable the BGP control plane for all VNIs using the configuration shown below.

```
admin@router# set protocol bgp evpn advertise-all-vni
admin@router# commit
Commit OK.
Save Done!
```

## Note



Only leaf switches that are VTEPs need this configuration. EVPN routes are still accepted from BGP peers as they reside in the global EVPN routing table, but are only made effective when the VNI corresponding to the received route is locally known.

## Sample Configuration Topology with EVPN Symmetric Routing

Figure 1 shows our sample topology for EVPN symmetric routing. We have two routers with two servers connected to each router.




Figure 1 EVPN Symmetric Routing

## Router Configuration

The configuration for Router 1 is shown below. The first part of the configuration deals with physical interfaces and assigning VLANs to these interfaces. Router 1 has three physical interfaces configured, two interfaces connecting server 1 and server 3.

Next we setup the layer 3 VLAN interfaces and configure IP addresses for these interfaces. The MTU value of layer 3 interfaces is set to 1450 to ensure there is enough space for the overlay VXLAN header. The switch will complain and refuse to commit the configuration if MTU size is not set to 1450.

Then we setup VXLAN VNIs, enable VLAN to VNI mapping and map VNIs to VLAN IDs. Finally we enable IP routing on the device and configure BGP related parameters such as router ID, BGP neighbor and enabling advertising all VNIs.

Follow the detailed configuration steps below.

### Router 1 Configuration

**Step 1.** Configure physical interfaces, VLAN interfaces and assign VLAN IDs and IP addresses.

```
admin@router1# set interface gigabit-ethernet te-1/1/1 family ethernet-switching native-vlan-id 100
admin@router1# set interface gigabit-ethernet te-1/1/10 family ethernet-switching native-vlan-id 1111
admin@router1# set interface gigabit-ethernet te-1/1/10 family ethernet-switching port-mode "trunk"
admin@router1# set interface gigabit-ethernet te-1/1/10 family ethernet-switching vlan members 2221
admin@router1# set interface gigabit-ethernet te-1/1/11 family ethernet-switching native-vlan-id 2222
admin@router1# set l3-interface loopback address 201.201.201.201 prefix-length 32
admin@router1# set l3-interface loopback address 201.201.201.88 prefix-length 32
admin@router1# set l3-interface loopback address 201::201 prefix-length 128
admin@router1# set l3-interface vlan-interface vlan100 mtu 1450
admin@router1# set l3-interface vlan-interface vlan100 address 100.1.1.201 prefix-length 24
admin@router1# set l3-interface vlan-interface vlan100 address 100::201 prefix-length 64
admin@router1# set l3-interface vlan-interface vlan2221 vrf "vrf1"
admin@router1# set l3-interface vlan-interface vlan2221 mtu 1450
admin@router1# set l3-interface vlan-interface vlan2221 address 21.1.1.201 prefix-length 24
admin@router1# set l3-interface vlan-interface vlan2222 vrf "vrf1"
admin@router1# set l3-interface vlan-interface vlan2222 mtu 1450
admin@router1# set l3-interface vlan-interface vlan2222 address 22.1.1.201 prefix-length 24
admin@router1# set l3-interface vlan-interface vlan1111 vrf "vrf1"
admin@router1# set l3-interface vlan-interface vlan1111 mtu 1450
admin@router1# set vlans vlan-id 100 l3-interface "vlan100"
admin@router1# set vlans vlan-id 1111 l3-interface "vlan1111"
admin@router1# set vlans vlan-id 2221 l3-interface "vlan2221"
admin@router1# set vlans vlan-id 2222 l3-interface "vlan2222"
```

**Step 2:** Configure VXLAN VNI and map VNI IDs to VLAN IDs. Also create an L3 VNI in vrf1.

```
admin@router1# set vxlans source-interface lo address 201.201.201.201
admin@router1# set vxlans vni 9999 vlan 1111
admin@router1# set vxlans vni 22221 vlan 2221
admin@router1# set vxlans vni 22222 vlan 2222
admin@router1# set vxlans vrf vrf1 l3-vni 9999
```

**Step 3:** Enable IP routing and configure VRF.

```
admin@R1# set ip routing enable true
admin@R1# set ip vrf vrf1
```

**Step 4:** Configure BGP related configuration.

```
admin@router1# set protocols bgp local-as 201
admin@router1# set protocols bgp router-id 201.201.201.201
admin@router1# set protocols bgp neighbor 100.1.1.134 remote-as "external"
admin@router1# set protocols bgp neighbor 100.1.1.134 update-source "100.1.1.201"
admin@router1# set protocols bgp neighbor 100.1.1.134 evpn activate
admin@router1# set protocols bgp ipv4-unicast network 201.201.201.201/32
admin@router1# set protocols bgp ipv4-unicast network 201.201.201.88/32
admin@router1# set protocols bgp evpn advertise-all-vni
admin@router1# set protocols bgp evpn advertise ipv4-unicast
admin@router1# set protocols bgp vrf vrf1 local-as 201
admin@router1# set protocols bgp vrf vrf1 router-id 201.201.201.201
admin@router1# set protocols bgp vrf vrf1 ipv4-unicast network 22.1.1.0/24
admin@router1# set protocols bgp vrf vrf1 evpn advertise ipv4-unicast
```

## Router 2 Configuration

Configuration of Router 2 is shown below. Router 2 also has 3 physical interfaces configured with two interfaces connecting server 2 and server 4. The configuration of Router 2 is almost identical to Router 1 other than the basic configurations like VXLAN VNI and interface IP addresses. The detailed configuration steps are shown below.

**Step 1.** Configure physical interfaces, VLAN interfaces and assign VLAN ID to physical interfaces and IP addresses.

```

admin@router2# set interface gigabit-ethernet ge-1/1/1 family ethernet-switching native-vlan-id 100
admin@router2# set interface gigabit-ethernet ge-1/1/2 disable true
admin@router2# set interface gigabit-ethernet ge-1/1/2 family ethernet-switching native-vlan-id 2221
admin@router2# set interface gigabit-ethernet ge-1/1/10 family ethernet-switching native-vlan-id 1111
admin@router2# set interface gigabit-ethernet ge-1/1/10 family ethernet-switching port-mode "trunk"
admin@router2# set interface gigabit-ethernet ge-1/1/10 family ethernet-switching vlan members 2221
admin@router2# set interface gigabit-ethernet ge-1/1/10 family ethernet-switching vlan members 2222
admin@router2# set interface gigabit-ethernet ge-1/1/11 family ethernet-switching native-vlan-id 2222
admin@router2# set interface gigabit-ethernet ge-1/1/11 family ethernet-switching port-mode "trunk"
admin@router2# set interface gigabit-ethernet ge-1/1/11 family ethernet-switching vlan members 3333
admin@router2# set l3-interface loopback address 134.134.134.134 prefix-length 32
admin@router2# set l3-interface vlan-interface vlan2221 vrf "vrf1"
admin@router2# set l3-interface vlan-interface vlan2221 mtu 1450
admin@router2# set l3-interface vlan-interface vlan2221 address 21.1.1.134 prefix-length 24
admin@router2# set l3-interface vlan-interface vlan100 mtu 1450
admin@router2# set l3-interface vlan-interface vlan100 address 100.1.1.134 prefix-length 24
admin@router2# set l3-interface vlan-interface vlan100 address 100::134 prefix-length 64
admin@router2# set l3-interface vlan-interface vlan3333 vrf "vrf1"
admin@router2# set l3-interface vlan-interface vlan3333 mtu 1450
admin@router2# set l3-interface vlan-interface vlan3333 address 33.1.1.134 prefix-length 24
admin@router2# set l3-interface vlan-interface vlan1111 vrf "vrf1"
admin@router2# set l3-interface vlan-interface vlan1111 mtu 1450
admin@router2# set vlans vlan-id 100 l3-interface "vlan100"
admin@router2# set vlans vlan-id 1111 l3-interface "vlan1111"
admin@router2# set vlans vlan-id 2221 l3-interface "vlan2221"
admin@router2# set vlans vlan-id 2222 l3-interface "vlan2222"
admin@router2# set vlans vlan-id 3333 l3-interface "vlan3333"

```

#### Step 2: Configure VXLAN VNI and map VNI IDs to VLAN IDs.

```

admin@router2# set vxlans source-interface lo address 134.134.134.134
admin@router2# set vxlans vni-mapvlan true
admin@router2# set vxlans vni 9999 vlan 1111
admin@router2# set vxlans vni 22221 vlan 2221
admin@router2# set vxlans vni 22222 vlan 2222
admin@router2# set vxlans vni 33333 vlan 3333

```

#### Step 3: Enable IP routing and configure VRF and hostname.

```

admin@router2# set ip routing enable true
admin@router2# set ip vrf vrf1
admin@router2# set evpn vrf vrf1 vni 9999
admin@router2# set vxlans vrf vrf1 l3-vni 9999

```

#### Step 4: Configure BGP related configuration

```

admin@router2# set protocols bgp local-as 134
admin@router2# set protocols bgp router-id 134.134.134.134
admin@router2# set protocols bgp neighbor 100.1.1.201 remote-as "external"
admin@router2# set protocols bgp neighbor 100.1.1.201 update-source "100.1.1.134"
admin@router2# set protocols bgp neighbor 100.1.1.201 evpn activate
admin@router2# set protocols bgp ipv4-unicast network 134.134.134.134/32
admin@router2# set protocols bgp evpn advertise-all-vni
admin@router2# set protocols bgp vrf vrf1 local-as 134
admin@router2# set protocols bgp vrf vrf1 router-id 134.134.134.134
admin@router2# set protocols bgp vrf vrf1 ipv4-unicast network 33.1.1.0/24
admin@router2# set protocols bgp vrf vrf1 evpn advertise ipv4-unicast

```

## EVPN MAC Learning Process

In Figure 2, when Host 1 is first plugged into R1, Host 1 will start sending ARP and other basic networking traffic like DHCP. When R1 receives a packet from Host 1 for the first time, it will record its MAC address in its local MAC address table. Also, R1 will advertise an EVPN Type-2 route to R2. The route includes the local EVPN instance of R1, the VTEP IP address, the Host 1 MAC address and the L2VNI.

Upon receiving the EVPN Type-2 route from R1, R2 learns the MAC address of Host 1. To accept this route, R2 needs to determine if the Import Route Target (IRT) configured on R2 matches the Export Route Target (ERT). RT is sent as the BGP Extended Community attribute. In this case the IRT and ERT match hence the route is accepted and the MAC address of Host 1 is learned.

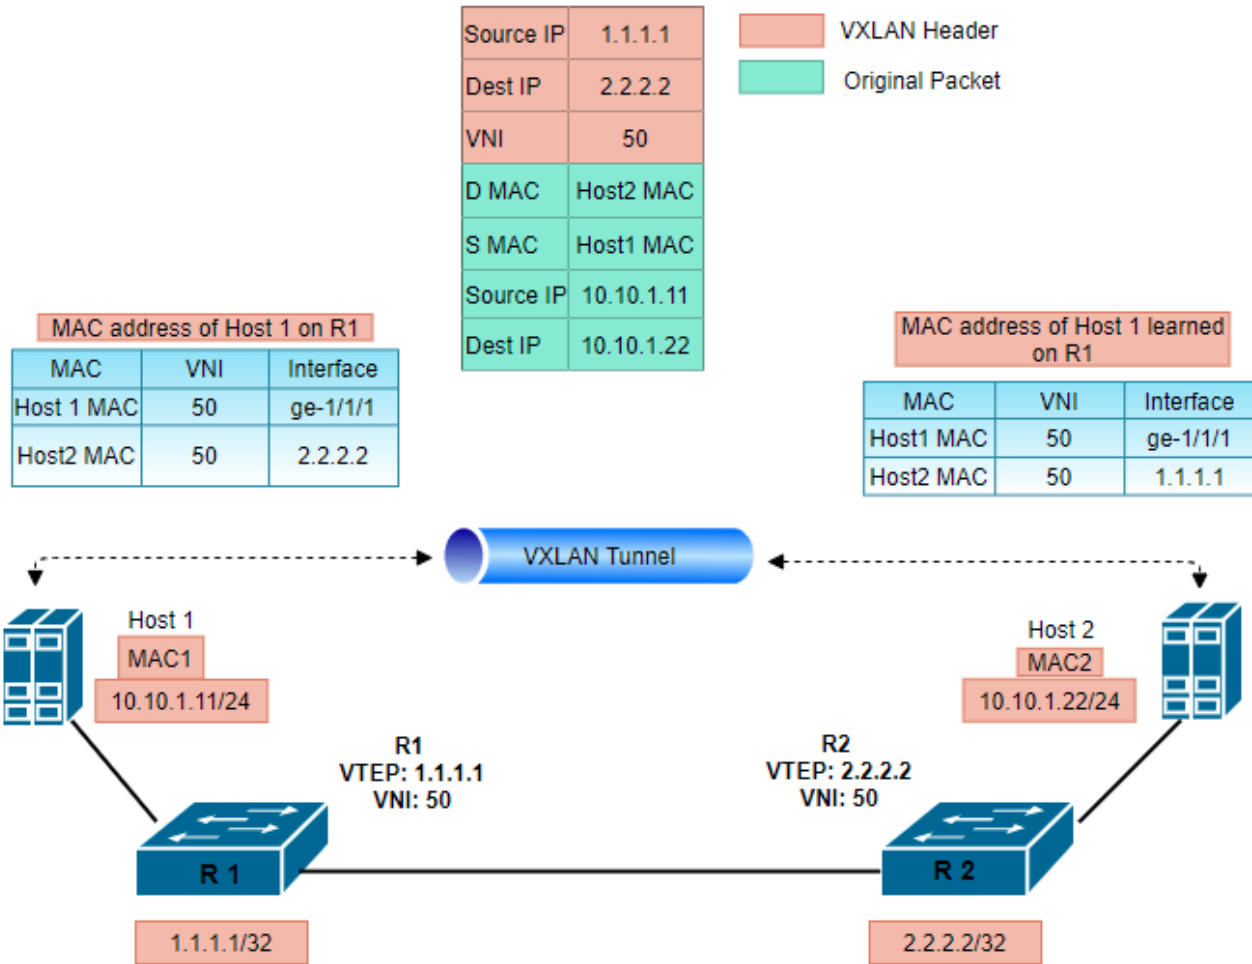


Figure 2. MAC Learning and Packet Forwarding

## Packet Forwarding Process

In the case of packet forwarding within the same subnet as depicted in Figure 2, both Host1 and Host2 belong to the same VNI. Host1 wants to send a packet to Host2.

1. If Host1 doesn't have the MAC address of Host2, the MAC address can be learned through the MAC learning process described in the section above. Assuming that Host1 does have the MAC address of Host2, Host 1 sends the packet to R1 destined for Host2.
2. R1 receives the packet and determines the VNI of Host1 based on the ingress interface configuration. R1 learned the Host2 MAC address and the outgoing interface is the VTEP 2.2.2.2 on R2.
3. R1 then encapsulates the original packet from Host1 with the VXLAN header and sends it out.
4. When the packet is received on R2, the outer VXLAN header is stripped off. R2 then searches its local MAC table and finds out the outgoing interface and delivers the original packet to Host2.

## Verifying Configuration

To check the BGP state and neighbor status on Router 2, we will run the `run show bgp neighbor` command.

```

admin@router2# run show bgp neighbor 100.1.1.201
BGP neighbor is 100.1.1.201, remote AS 201, local AS 134, external link
Hostname: 51.201
Member of peer-group fabric for session parameters
  BGP version 4, remote router ID 201.201.201.201, local router ID 134.134.134.134
  BGP state = Established, up for 01:22:43

```

To check the BGP EVPN routes on Router 2, we will run the command **run show bgp evpn route** on Router 2 as shown below. In our sample configuration we haven't manually configured route target (RD), so the RDs we see in the command output below are automatically configured by the system.

```

admin@router2# run show bgp evpn route
BGP table version is 5, local router ID is 134.134.134.134
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN type-1 prefix: [1]:[ESI]:[EthTag]:[IPlen]:[VTEP-IP]
EVPN type-2 prefix: [2]:[EthTag]:[MAClen]:[MAC]:[IPlen]:[IP]
EVPN type-3 prefix: [3]:[EthTag]:[IPlen]:[OrigIP]
EVPN type-4 prefix: [4]:[ESI]:[IPlen]:[OrigIP]
EVPN type-5 prefix: [5]:[EthTag]:[IPlen]:[IP]

      Network          Next Hop          Metric LocPrf Weight Path
      Extended Community
Route Distinguisher: 134.134.134.134:2
*> [5]:[0]:[24]:[33.1.1.0]
      134.134.134.134          0          32768 i
      ET:8 RT:134:9999 Rmac:04:f8:f8:20:6c:7b
Route Distinguisher: 134.134.134.134:3
*> [3]:[0]:[32]:[134.134.134.134]
      134.134.134.134          32768 i
      ET:8 RT:134:22221
Route Distinguisher: 134.134.134.134:4
*> [3]:[0]:[32]:[134.134.134.134]
      134.134.134.134          32768 i
      ET:8 RT:134:22222
Route Distinguisher: 134.134.134.134:5
*> [3]:[0]:[32]:[134.134.134.134]
      134.134.134.134          32768 i
      ET:8 RT:134:33333
Route Distinguisher: 201.201.201.201:2
*> [5]:[0]:[24]:[22.1.1.0]
      201.201.201.201          0          0 201 i
      RT:201:9999 ET:8 Rmac:c4:39:3a:fb:be:d9
Route Distinguisher: 201.201.201.201:3
*> [3]:[0]:[32]:[201.201.201.201]
      201.201.201.201          0 201 i
      RT:201:22221 ET:8
Route Distinguisher: 201.201.201.201:4
*> [3]:[0]:[32]:[201.201.201.201]
      201.201.201.201          0 201 i
      RT:201:22222 ET:8

```

To verify the VXLAN tunnel information, run the command **run show vxlan tunnel** on Router 2.

```

admin@router2# run show vxlan tunnel
Total number of tunnels: 3

VNI 9999, Encap:service-vlan-delete, Decap:service-vlan-add-replace
src addr:134.134.134.134, dst addr:201.201.201.201, state:UP
traffic type:unicast
nexthops:100.1.1.201
output ports:ge-1/1/1

VNI 22221, Encap:service-vlan-delete, Decap:service-vlan-add-replace
src addr:134.134.134.134, dst addr:201.201.201.201, state:UP
traffic type:unicast
nexthops:100.1.1.201
output ports:ge-1/1/1

VNI 22222, Encap:service-vlan-delete, Decap:service-vlan-add-replace
src addr:134.134.134.134, dst addr:201.201.201.201, state:UP
traffic type:unicast
nexthops:100.1.1.201
output ports:ge-1/1/1

```

Run the command **run show vxlan evpn rmac** to display the Router-MAC (rmac) and other VXLAN parameters. Note that to configure rmac for a layer 3 VLAN interface, run the command **set l3-interface vlan-interface <interface-name>router-mac <router-mac>**

```

admin@router2# run show vxlan evpn rmac

```

L3 VNI	Interface	SVI-Interface	VTEP	RMAC	Flags
9999	vxlan9999	vlan1111	201.201.201.201	c4:39:3a:fb:be:d9	0x16

To display the VXLAN EVPN routes, run the command **run show vxlan evpn route**.

```

admin@router2# run show vxlan evpn route

```

VRF-ID	ROUTE	NextHop	VNI	Interface
1	22.1.1.0/24	201.201.201.201	9999	vlan100

To check the routing table of a device, run the command **run show route vrf vrf1**. We can see that there is a BGP route for subnet 22.1.1.0/24 via 201.201.201.201 to reach server 3.

```

admin@router2# run show route vrf vrf1
show ip route vrf vrf1
=====
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, E - EIGRP, N - NHRP,
       T - Table, v - VNC, V - VNC-Direct, A - Babel, D - SHARP,
       F - PBR, f - OpenFabric,
       > - selected route, * - FIB route, q - queued route, r - rejected route

VRF vrf1:
K>* 0.0.0.0/0 [255/8192] unreachable (ICMP unreachable), 00:32:26
C>* 21.1.1.0/24 is directly connected, vlan2221, 00:31:37
B>* 22.1.1.0/24 [20/0] via 201.201.201.201, vlan1111 onlink, weight 1, 00:31:24 <----- Route for
server3
C>* 33.1.1.0/24 is directly connected, vlan3333, 00:31:37

```