# An EVPN routing cannot learn problem handling experience case

Switches 唐勋 2021-03-26 Published

## Network Topology

In the EVPN networking, Border is connected to the external network through VPC1 (VRFA). The VM belongs to VPC2 (VRFB) and hangs under the leaf. Border passes the routes learned from the external network to leaf through EVPN, and connects the two VPCs (using the BGP RT attribute to enable routes to learn from each other).

## Problem Description

The remote Border device advertises the prefix route 110.1.1.1.0/24 through the EVPN control plane. T he local leaf checks the system routing table of the VPN VRFB, but there is no such route.

## Process Analysis

1. Check whether the EVPN neighbor status between the device and the RR is normal. Established is the normal state.

```
display bgp   instance SDN_INSTANCE_BGP peer l2vpn evpn

  BGP local router ID: 2.2.2.6

  Local AS number: 65205

  Total number of peers: 2                                 Peers in established state:
2

   * - Dynamically created peer

   Peer               AS   MsgRcvd   MsgSent OutQ PrefRcv Up/Down   State


   2.2.2.4       65205         1468         1729       0             6 21:48:01 Establ
ished
   2.2.2.5       65205         1407         1461       0             6 21:48:05 Establ
ished
```

2. Prefix routing is delivered through TYPE 5 routing messages. The sending and receiving of TYPE 5 prefix routing is controlled by the RT configured in the address-family ipv4 view under the ip vpn-instance instance. Second, you need to check whether the RT attributes configured at both ends meet the requirements for BGP route import (the RT attribute carried in the route must have at least one value belonging to the import-extcommunity list configured on the receiving end).

```
ip vpn-instance VRFB
  route-distinguisher 1:24001
  #
  address-family ipv4
   export route-policy SDN_POLICY_EXPORT_IPV4_VRFB
   vpn-target 0:24001 1:24001 0:25001 import-extcommunity
   vpn-target 1:24001 export-extcommunity
```

3. Check the EVPN TYPE5 route in BGP to check whether there is a prefix route advertised by the remote end. It can be seen from the following information that the remote prefix route can be received, and the status in the BGP routing table is also valid, and the optimal route is also calculated.

```
[LA1]display   bgp   instance SDN_INSTANCE_BGP l2vpn   evpn   route-type ip-prefix
  BGP local router ID is 2.2.2.6
  Status codes: * - valid, > - best, d - dampened, h - history,
                           s - suppressed, S - stale, i - internal, e - external
                           Origin: i - IGP, e - EGP, ? - incomplete


  Total number of routes from all PEs: 9


  Route distinguisher: 4:25001
  Total number of routes: 4


        Network               NextHop           MED           LocPrf         PrefVal
Path/Ogn


* >i [5][0][24][110.1.1.0]/80
                                    2.2.2.1               0                  10
0             0             i
*   i                           2.2.2.1               0                  10
0             0             i
```

4. Check whether the prefix route is generated normally in the evpn routing-table table, and it is found that the above-mentioned prefix route is not generated. It can be seen from this that the problem occurs in the process of adding the route optimized by BGP to the evpn routing-table table.

```
[LA1]display evpn routing-table vpn-instance VRFB

VPN instance:VRFB          Local L3VNI:24001
IP address                 Next hop              Outgoing interface       NibID

10.1.2.48                  2.2.2.8               Vsi-interface24001       0x1800000
0
```

Resolved after adding the relevant configuration of the L3VNI interface.

5. View the detailed information of the above prefix routing. It is found that the route attributes carried in the following routes are normal, such as extended community attributes such as RT and L3VNI.

```
[LA1]display bgp instance SDN_INSTANCE_BGP l2vpn evpn route-distinguisher 4:25001 ro
ute-type ip-prefix [5][0][24][110.1.1.0] 80


  BGP local router ID: 2.2.2.6

  Local AS number: 65205


  Route distinguisher: 4:25001

  Total number of routes: 2

  Paths:     2 available, 1 best


  BGP routing table information of [5][0][24][110.1.1.0]/80:

  From                       : 2.2.2.4 (2.2.2.4)

  Rely nexthop       : 3.3.0.33

  Original nexthop: 2.2.2.1

  OutLabel               : NULL

  Ext-Community      :   <RT: 0:25001>, <RT: 1:25001>, <Encapsulation Type: VXLAN>, <Router's MAC: 84
d9-31f5-3908>

  AS-path                 : (null)

  Origin                   : igp

  Attribute value : MED 0, localpref 100, pref-val 0

  State                     : valid, internal, best

  Originator           : 2.2.2.1

  Cluster list       : 2.2.2.4

  IP precedence     : N/A

  QoS local ID       : N/A

  Traffic index     : N/A

  EVPN route type : IP prefix advertisement route

  ESI                         : 00:00:00:00:00:00:00:00:00:00

  Ethernet tag ID : 0

  IP prefix             : 110.1.1.0/24

  Gateway address : 0.0.0.0

  MPLS label           : 25001
```

6. To generate a route in the routing table, the next hop of this route must be reachable. The EVPN networking is quite special. In the implementation of symmetric IRB forwarding, the next hop of the remote route is the loopback port address of the remote VTEP. The device needs to generate the arp entry of the next hop, and the MAC address in the arp entry is The Router's MAC: 84d9-31f5-3908 value in the extended community attribute carried in the above detailed routing information. At the same time, the outgoing interface is the VSI interface of the L3VNI carried in the route. Check whether an ARP entry for the next hop is formed on the device. It can be seen from the following information that the device does not generate the next hop arp information, which eventually causes the routing to fail to take effect.

```
[LA1]display arp 2.2.2.1
```