Confidentiality Controlled

Typical Features of RG-S6220-H Series Data Center Switches

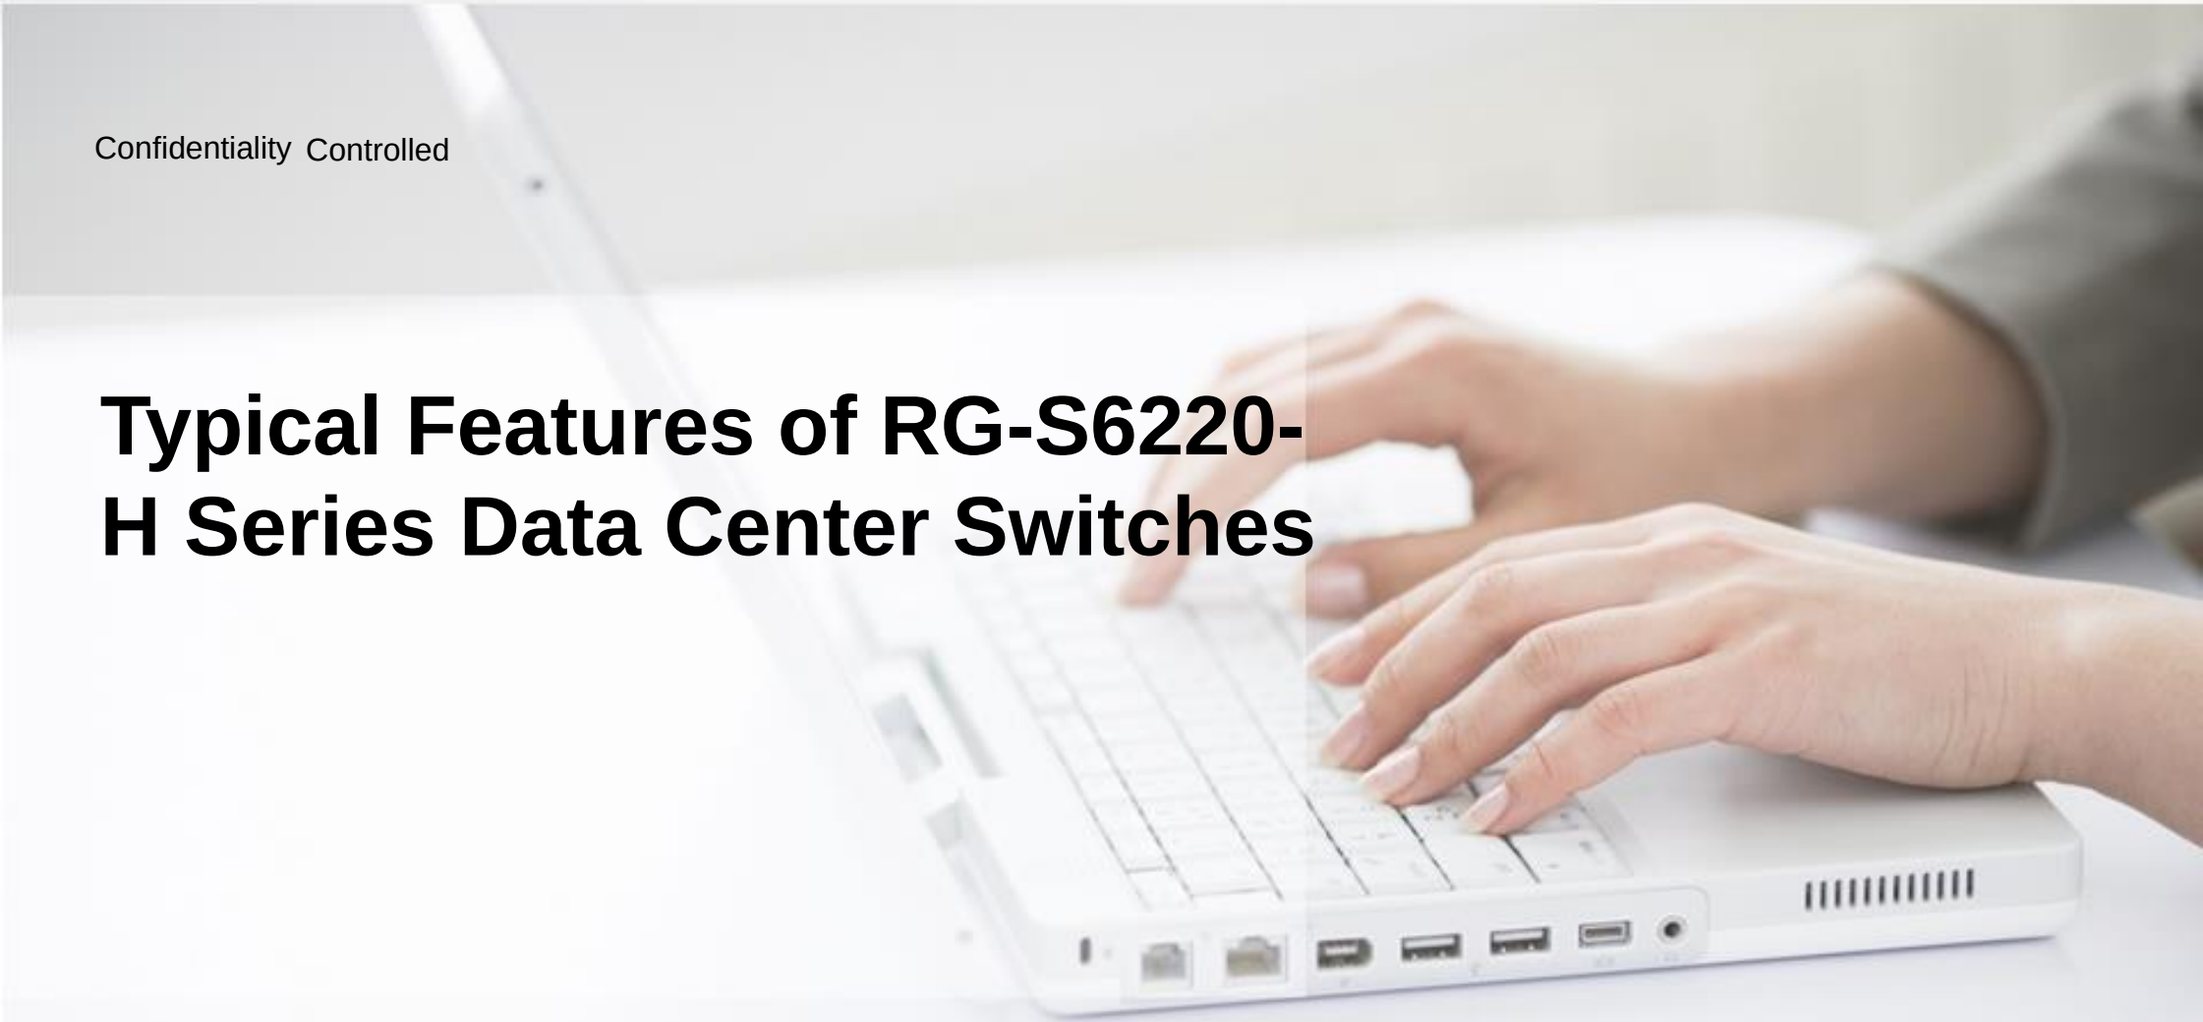

Contents

- **Elastic Hashing**
- **VXLAN**
- **Optimized 40G Interface Breakout Display**

Elastic Hashing

- In a data center, a load balancing cluster generally connects to top of rack (TOR) switches using the Equal Cost Multipath (ECMP) protocol. TOR switches use ECMP to evenly distribute data flows to members in the load balancing cluster.

Problem: If a server fails, the traditional ECMP mechanism will re-distribute all data flows. For example, a flow previously distributed to server 1 may be distributed to server 2, but server 2 does not have historical information about the session, leading to interruption of the session.

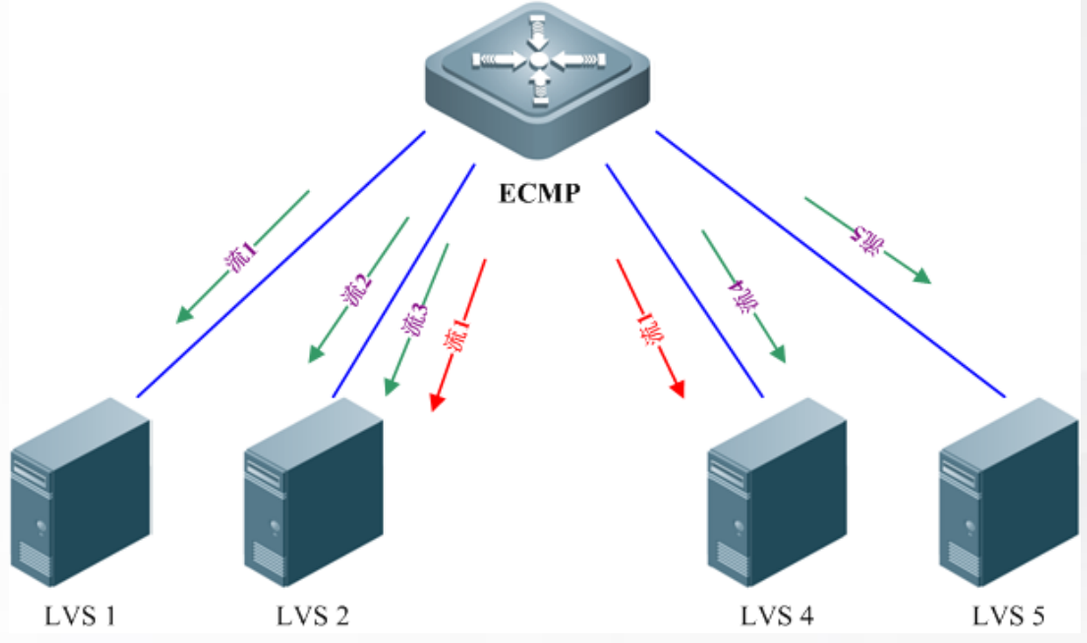

Elastic Hashing

- Elastic hashing (ECMP cluster) overcomes the problem of traffic redistribution caused by changes of ECMP paths. If the number of ECMP paths decreases, the ECMP cluster feature only distributes traffic originally on failed links to active links and retains traffic on active links. If the number of ECMP paths increases, this feature distributes some of traffic to the new active links.

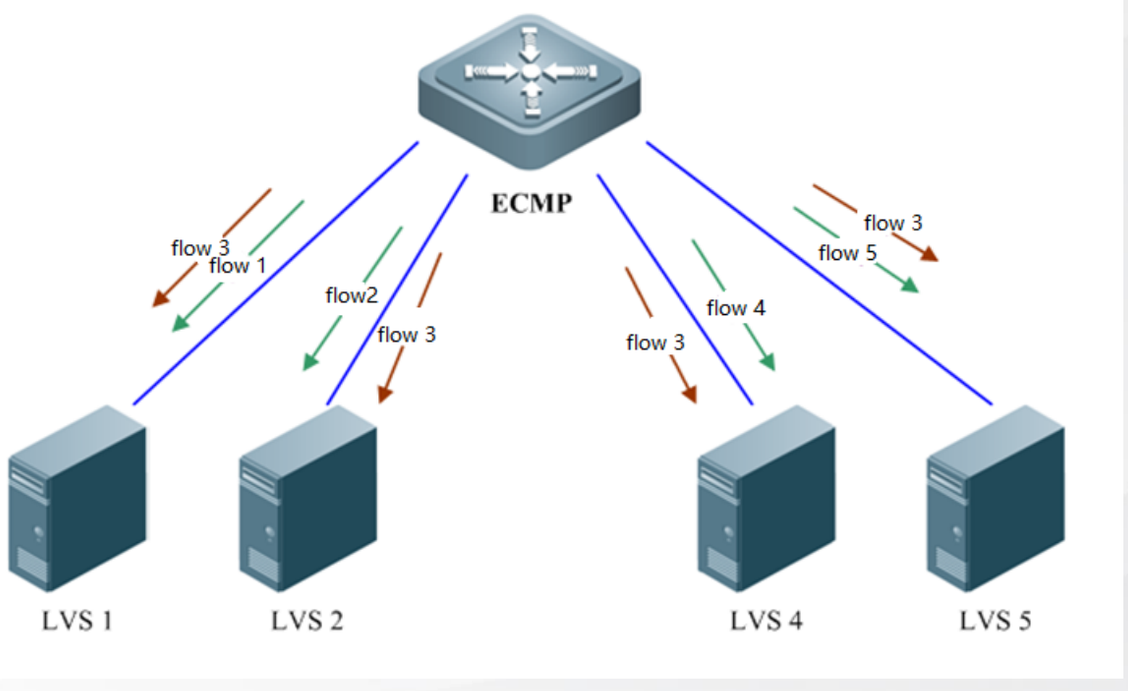

Elastic Hashing

- Typical scenario
  - A TOR switch connects to an LVS load balancing scheduler cluster through ECMP.

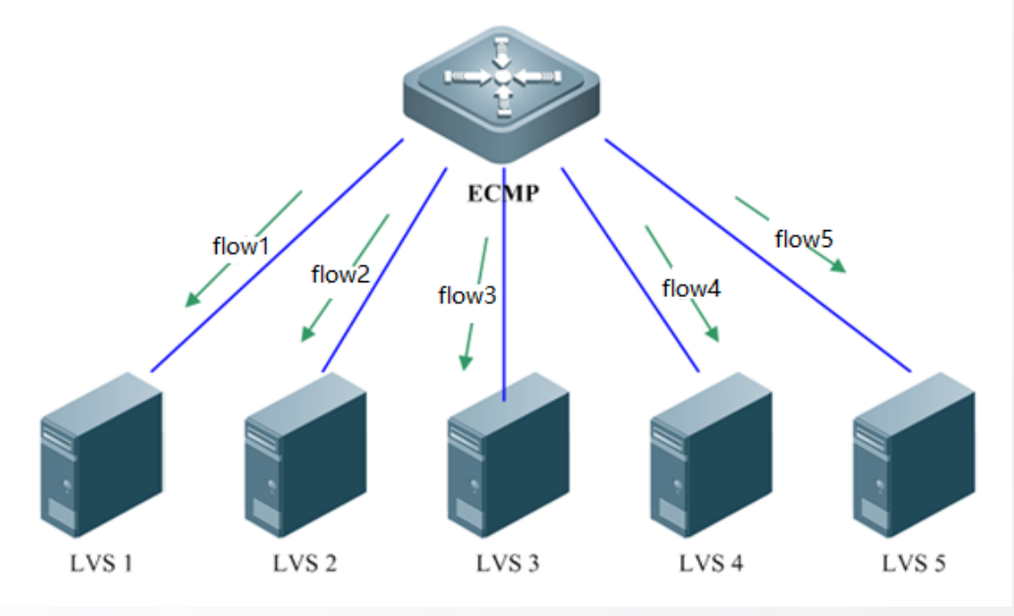

- Function enabled

  Ruijie(config)#ecmp cluster enable

- Function disabled

  Ruijie(config)#no ecmp cluster enable

Contents

- Elastic Hashing
- VXLAN
- Optimized 40G Interface Breakout Display

EVPN VXLAN

Virtual eXtensible Local Area Network (VXLAN) is an overlay technology that establishes an overlay virtual L2 network on an L3 network by encapsulating L2 packets in UDP tunnel packets.

- ✓ Uses mature, stable IP networks as bearer networks.
- ✓ Allows for free migration of VMs between servers across L3 network without being aware of the bearer network.
- ✓ Uses UDP encapsulation so that intermediate devices can implement load balancing of VXLAN packets based on IP 5-tuple information in outer tags, without awareness of VXLAN.
- ✓ Uses a new subnet identifier, 24-bit VXLAN network identifier (VNI), to define a logical L2 network.
- ✓ For more information about VXLAN fundamentals, refer to the *VXLAN Technology White Paper*.

EVPN VXLAN

- RFC7348 only defines the VXLAN data plane and does not define the VXLAN control plane. VXLAN devices learn remote virtual tunnel endpoint (VTEP) and host information depending on flooding on the data plane. Flooding becomes an obstacle to large-scale VXLAN deployment.

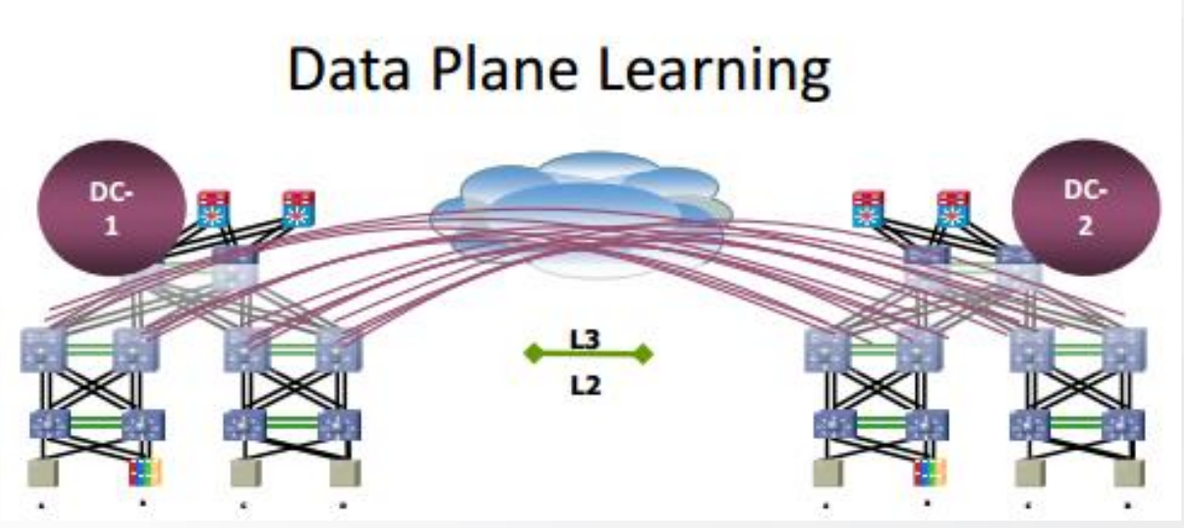

- The VXLAN data plane transmits flooding traffic in IP multicast packets. Therefore, multicast protocols need to be deployed on the underlay network. Defects of multicast protocols, such as complex management, protocol cost, and number of multicast groups, become an obstacle to VXLAN deployment.

- Generally, network operators do not open multicast services on their wide area networks to common users. Therefore, the current VXLAN technology cannot support VXLAN deployment across networks of different operators.

EVPN VXLAN

Multiprotocol Border Gateway Protocol Ethernet Virtual Private Network (MP-BGP EVPN) is considered a factual standard for the VXLAN control plane.

EVPN implements VTEP discovery and synchronization of host reachability information, thereby reducing flooding traffic on the network and avoiding dependency on multicast services on the underlay network. This feature makes large-scale VXLAN deployment possible.

S6220-H is the first series of Ruijie data center switches supporting the EVPN VXLAN solution. VXLAN can be deployed in centralized or distributed mode, depending on whether the VXLAN gateways for VMs are deployed on core or access devices. The distributed VXLAN deployment is used as an example to describe the configuration procedure later.

EVPN VXLAN

Scenario

Distributed VXLAN deployment is generally applicable to medium or large-sized data centers. After TOR switches are replaced by VXLAN routing-capable switches (such as S6220-H), an overlay VXLAN virtual network can be deployed on the existing L3 network. This deployment enables VMs to migrate flexibly among physical servers.

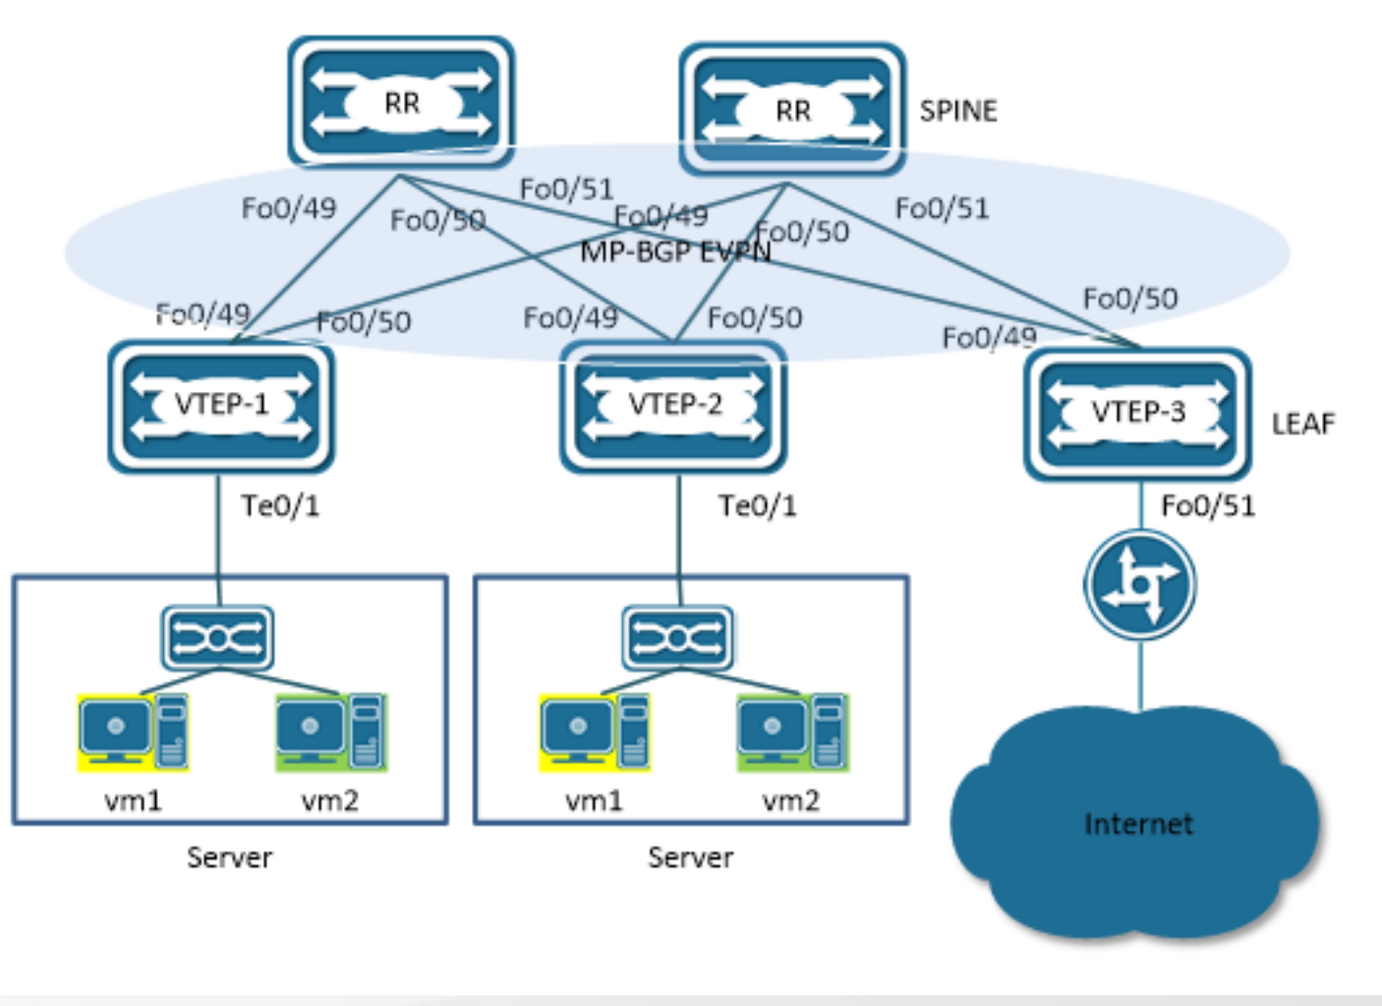

Configuration procedure

1. Change the working mode of VTEP-1/2/3 to EVPN-VXLAN.
2. Configure unicast routing to ensure unicast route reachability on the underlay network.
3. Configure VTEP functions and enable the BGP EVPN address family.
4. Configure VXLAN/VLAN information for users, configure VXLAN gateways, and bring hosts online.
5. Configure the border leaf node to enable access to external networks from the hosts.

EVPN VXLAN

■ **Step 1** Change the working mode of VTEPs.

1. Change the UFT mode of VTEP-1 to VXLAN.

VTEP-1(config)#switch-mode vxlan slot 0

Please save current config and restart your device!

VTEP-1(config)#end

VTEP-1#write ←Save the configuration as prompted.

2. Set the VXLAN working mode to EVPN.

VTEP-1# vxlan convert mode evpn

Convert VXLAN Mode to EVPN and Reload System(make sure configuration has been saved)? (y/N)

←At the prompt, enter **y**, and the switch restarts automatically.

3. Repeat the preceding steps on the other VTEPs.

(Detailed steps not provided here.)

EVPN VXLAN

■ **Step 2** Configure unicast routing.

1. Configure IP addresses for physical interfaces.

VTEP-1(config)# interface fortyGigabitEthernet 0/ 49

VTEP-1(config-if-FortyGigabitEthernet 0/49)# no switchport

VTEP-1(config-if-FortyGigabitEthernet 0/49)# ip address 11.1.1.2 255.255.255.0

...

2. Configure a loopback interface address to identify the local device.

VTEP-1(config)# interface loopback 1

VTEP-1(config)# ip address 1.1.1.1 255.255.255.0

3. Configure a dynamic routing protocol, for example, OSPF.

VTEP-1(config)# router bgp 100

VTEP-1(config)# network 1.1.1.1 0.0.0.0 area 0

VTEP-1(config)# network 11.1.1.0 0.0.0.255 area 0

4. Configure L3 interfaces and loopback interface addresses on VTEP

(Detailed steps not provided here.)

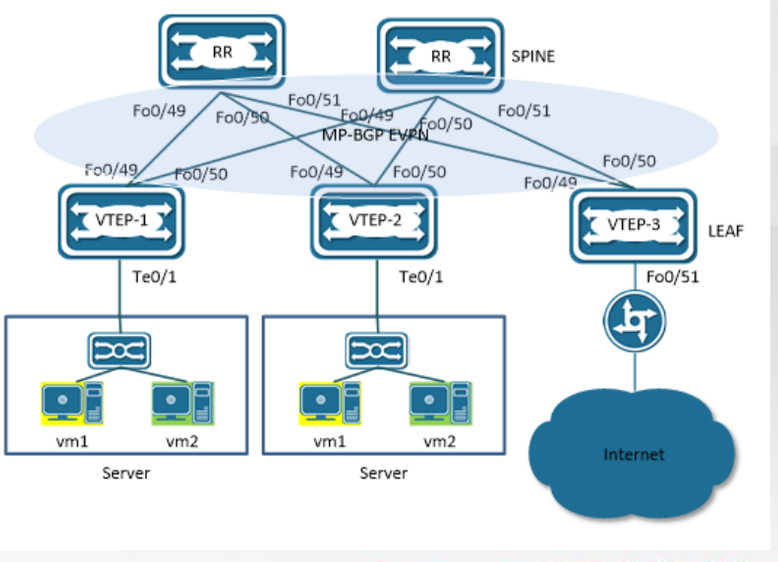

EVPN VXLAN

■ **Step 3** Configure VTEPs and BGP EVPN.

1. Configure VTEPs.

VTEP-1(config)# vtep

VTEP-1(config-vtep)# source loopback 1

VTEP-1(config-vtep)# fabric anycast-gateway-mac 0000.1111.2222 ←Configure an anycast virtual MAC address ，(optional).

VTEP-1(config-vtep)# arp suppress enable ←Enable ARP suppression，(optional).

2. Enable BGP.

VTEP-1(config)# router bgp 100

VTEP-1(config-router)# neighbor 3.3.3.3 remote-as 100 ←Specify the route reflector (RR) as the BGP neighbor.

VTEP-1(config-router)# neighbor 3.3.3.3 update-source Loopback 1

VTEP-1(config-router)# address-family l2vpn evpn ←Enable the EVPN address family.

VTEP-1(config-router-af)# neighbor 3.3.3.3 activate

3. Configure the BGP RR.

RR-1(config)# router bgp 100

RR-1(config-router)# neighbor 1.1.1.1 remote-as 100

RR-1(config-router)# neighbor 1.1.1.1 update-source Loopback 0

RR-1(config-router)# address-family l2vpn evpn

RR-1(config-router-af)# neighbor 1.1.1.1 route-reflector-client ←Specify 1.1.1.1 as the RR client.

EVPN VXLAN

■ **Step 4** Configure VXLAN/VLAN information for users.

1. Specify VLANs for users.

```
VTEP-1(config)# vlan 10
VTEP-1(config)# interface tenGigabitEthernet 0/1
VTEP-1(config-if-TenGigabitEthernet  0/1)# switch access vlan 10
```

2. Configure VLAN to VXLAN mapping.

```
VTEP-1(config)# vxlan 100
VTEP-1(config-vxlan)# extend-vlan 10
VTEP-1(configr)# ip vrf tenant.1 ←Create tenant VRFs.
VTEP-1(configr)# interface overlayRouter 100
VTEP-1(config-if-OverlayRouter 100)# ip address 172.18.159.1 255.255.255.0
VTEP-1(config-if-OverlayRouter 100)# ip vrf forwarding tenant.1 ←Specify the VRF to which the gateway belongs.
VTEP-1(config-if-OverlayRouter 100)# anycast-gateway ←Configure the gateway as an anycast gateway.
VTEP-1(config-vxlan)# router-interface overlayrouter 100 ←Specify the gateway for VXLAN 100.
```

3. Configure EVI.

```
VTEP-1(config)# evpn
VTEP-1(config-evpn)# vni 100 ←Configure one-to-one mapping between the EVI， and VXLAN instance ID.
VTEP-1(config-evpn-vni)# rd auto
VTEP-1(config-evpn-vni)# route-target both auto
```

4. Repeat steps 2 and 3 for other VXLAN instances and complete the configurations on other VTEPs.

EVPN VXLAN

■ **Step 5** Configure the border leaf node.

1. Configure a route to the public network.

Configure a static route in the default VRF, as shown in the red text in the figure. Or, use a routing protocol to advertise the route to the public network to the underlay network.

2. Configure an outbound interface for traffic from tenants.

```
Border(config)# int Fo0/51
Border (config-if-FortyGigabitEthernet 0/51) # switchport mode trunk
```
←Configure the interface connected to the public network as a trunk interface.
```
Border(config)# int overlayrouter 100
Border(config-if-overlayrouter 100)# ip vrf forwarding tenant.1
Border(config-if-overlayrouter 100)# ip addr 172.18.159.1/24
Border(config)# int overlayrouter 1001
Border(config-if-overlayrouter 100)# ip vrf forwarding tenant.1
Border(config-if-overlayrouter 100)# ip addr 101.1.1.1/24
Border(config-if-overlayrouter 100)# anycast-gateway
Border(config)# vxlan 100
Border(config-vxlan)# router-interface overlayrouter 100
Border(config)# vxlan 1001
```
←Configure VNI 1001 for access to the public network.
```
Border(config-vxlan)# extend-vlan 1001
```
←Assign VLAN 1001 for access to tenant 1 from external routers.
```
Border(config-vxlan)# router-interface overlayrouter 1001
```

3. Use the routing protocol to advertise intranet routes to the external routers. (Detailed steps not provided here.)

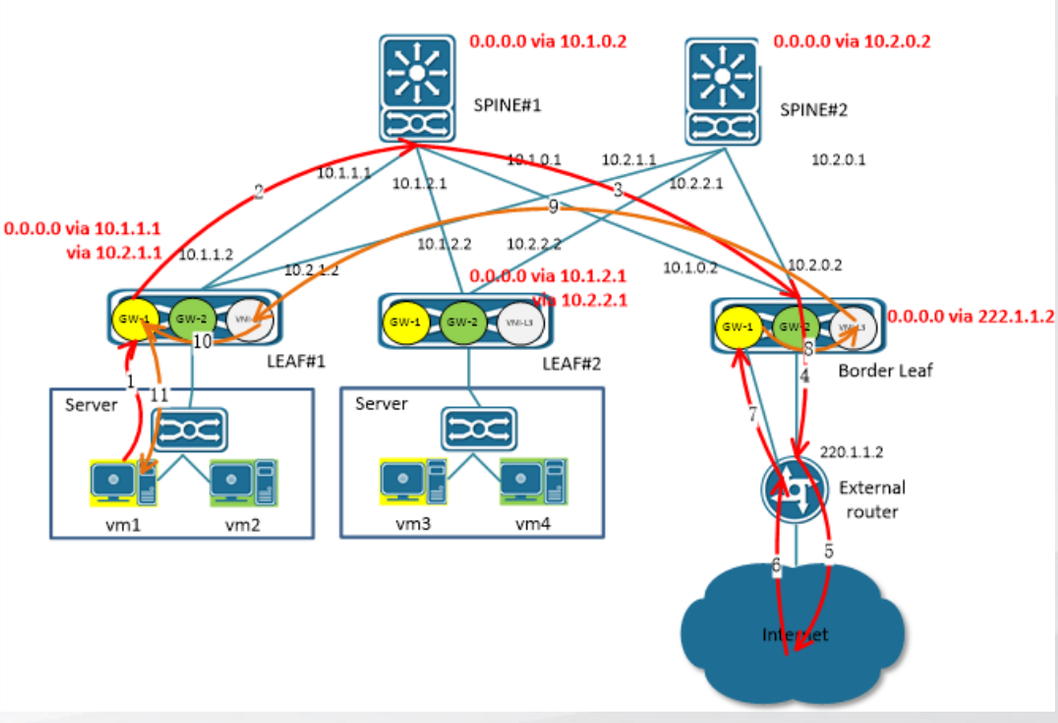

Note: This is a workaround solution. The formal solution requires support for VXLAN routing and is not provided now.

Contents

- **Elastic Hashing**
- **VXLAN**
- **Optimized 40G Interface Breakout Display**

Optimized 40G Interface Breakout Display

Original 40G Interface Number	Derived 10G Interface Number (Optimized)	Derived 10G Interface Number (Previous)
interface FortyGigabitEthernet 0/49	interface FortyGigabitEthernet 0/49:1 interface FortyGigabitEthernet 0/49:2 interface FortyGigabitEthernet 0/49:3 interface FortyGigabitEthernet 0/49:4	interface TenGigabitEthernet 0/55 interface TenGigabitEthernet 0/56 interface TenGigabitEthernet 0/57 interface TenGigabitEthernet 0/58
.....	.....	
interface FortyGigabitEthernet 0/54	interface FortyGigabitEthernet 0/54:1 interface FortyGigabitEthernet 0/54:2 interface FortyGigabitEthernet 0/54:3 interface FortyGigabitEthernet 0/54:4	interface TenGigabitEthernet 0/75 interface TenGigabitEthernet 0/76 interface TenGigabitEthernet 0/77 interface TenGigabitEthernet 0/78

For example, an S6220-48XS6QXS-H switch has 48 10G optical interfaces and 6 40G interfaces. If any 40G interface is divided into 10G interfaces, you can determine which 40G interface has been divided directly from numbers of the 10G interfaces.