

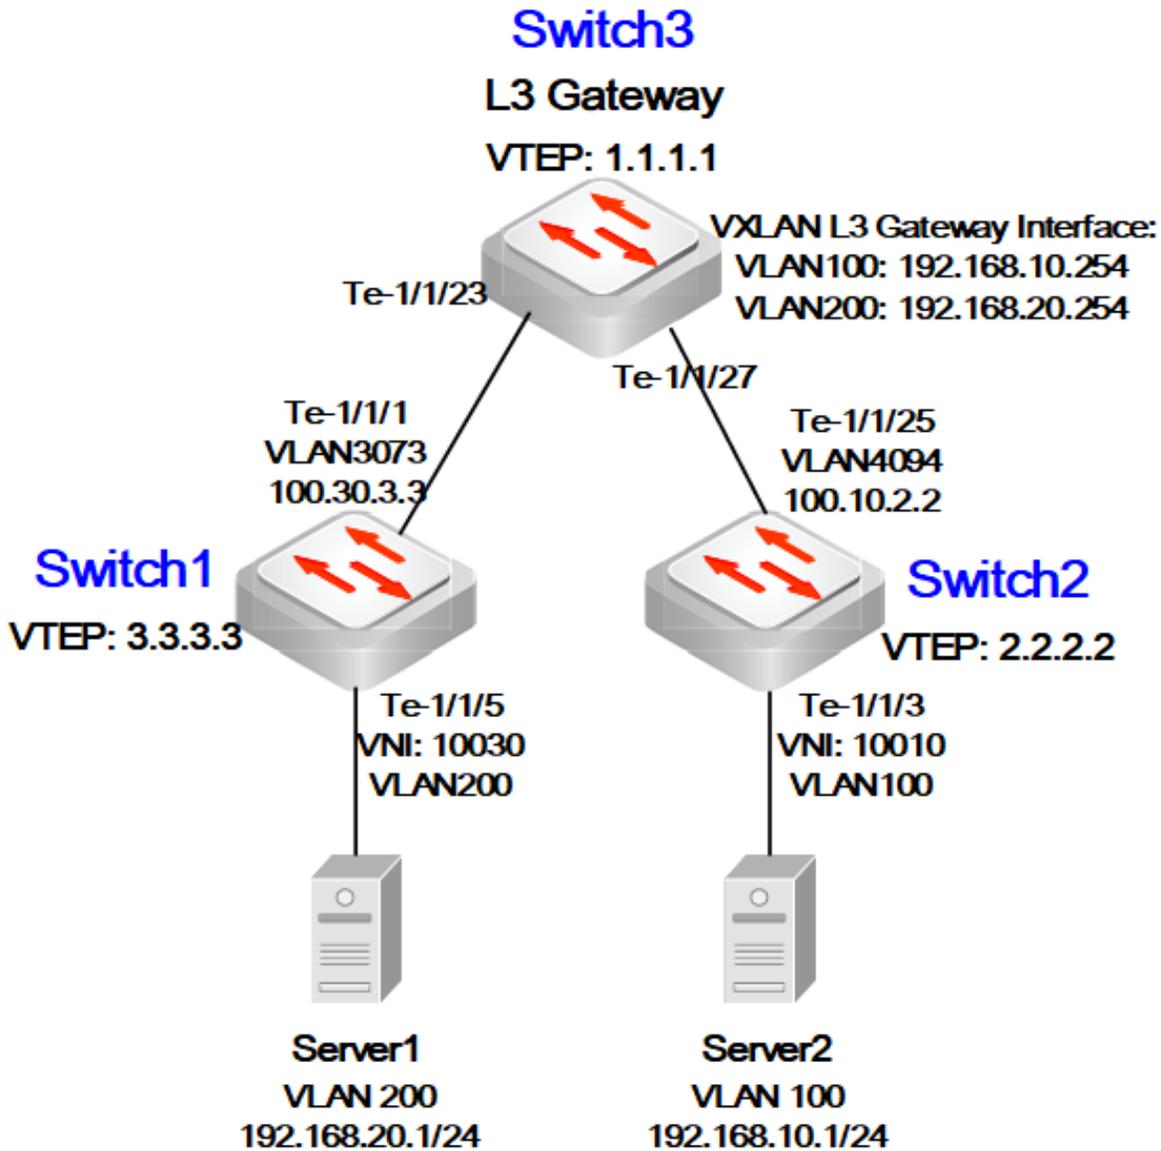
Cross-Subnet Packet Forwarding Process

Cross-subnet packet forwarding between VXLAN and VXLAN subnet or between VXLAN and non-VXLAN subnet can be implemented through a Layer 3 gateway.

Currently, PICOS only supports centralized gateway scenario. A centralized gateway is a Layer 3 gateway deployed centrally on a single device, where all cross-subnet traffic is forwarded through the Layer 3 gateway to achieve centralized management of traffic.

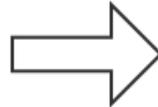
In the centralized gateway scenario, **Figure 1** shows the topology of cross-subnet packet forwarding and packet encapsulation.

Figure 1 Topology of cross-subnet packet forwarding and packet encapsulation



DMAC	NextHopMac1
SMAC	Switch1
SIP	3.3.3.3
DIP	1.1.1.1
VNI	10030
DMAC	L3 Gateway
SMAC	Server1
SIP	192.168.10.1
DIP	192.168.20.1
Payload	

VXLAN packet encapsulated by Switch1



DMAC	NextHopMac2
SMAC	Switch3
SIP	1.1.1.1
DIP	2.2.2.2
VNI	10010
DMAC	Server2
SMAC	L3 Gateway
SIP	192.168.10.1
DIP	192.168.20.1
Payload	

VXLAN packet encapsulated by Switch3



Cross-subnet IP packets from Server1



DMAC	L3 Gateway
SMAC	Server1
SIP	192.168.10.1
DIP	192.168.20.1
Payload	

IP packet decapsulated by Switch2



DMAC	Server2
SMAC	L3 Gateway
SIP	192.168.10.1
DIP	192.168.20.1
Payload	

The packet forwarding process is as follows:

1. Switch 1 receives the message from Server 1, obtains the corresponding layer 2 broadcast domain according to the access port and VLAN information in the message, and finds out the interface and encapsulation information in the corresponding Layer 2 broadcast domain.
2. VTEP on Switch 1 performs VXLAN message encapsulation based on the obtained egress interface and VNI information, and forwards it to Switch 3.
3. After receiving the VXLAN packet, Switch 3 decapsulates it and finds that the destination MAC in the inner packet is the MAC address of VLAN100, which is the VXLAN Layer 3 gateway interface. Switch 3 then determines that Layer 3 forwarding is needed.
4. Switch 3 strips off the Ethernet encapsulation of the inner packet, obtains the destination IP. Then Switch 3 looks up the routing table according to the destination IP to find the next hop address. Then Switch 3 looks up the ARP table entry according to the next hop address to obtain the destination MAC, VXLAN tunnel egress interface and VNI information.
5. Switch 3 re-encapsulates the VXLAN packet and forwards it to Switch 2, in which the source MAC in the Ethernet header of the inner packet is the MAC address of the VXLAN Layer 3 gateway interface VLAN 200.
6. After receiving the VXLAN packet on Switch 2, VTEP determines the validity of VXLAN packet based on UDP destination port number, source /destination IP address and VNI.

According to the VNI, the corresponding layer-2 broadcast domain is obtained. Switch 2 then obtains the interface and encapsulation information in the corresponding Layer 2 broadcast domain and performs VXLAN decapsulation on the packet.

7. Switch 2 forwards the message to Server 2 based on the egress interface and encapsulation information.

The process of message forwarding from Server 2 to Server 1 is similar and will not be repeated here.