



Layer-2 Multicast Forwarding

- [Multicast, on page 1](#)

Multicast

Overview

Layer-2 multicast traffic is typically transported in the overlay. In the VXLAN BGP EVPN fabric, Layer-2 multicast refers to the use of IP multicast technology for transporting multicast traffic in a Layer-2 switching environment.

Layer-2 multicast follows the same semantic as general Broadcast, Unknown Unicast, and Multicast (BUM) traffic. On Cisco Nexus 7000 Series and Cisco Nexus 9000 Series switches, IGMP snooping is used to facilitate Layer-2 multicast forwarding in the fabric.

This chapter only explains Layer-2 multicast. Multicast routing, used for transporting Layer-3 multicast traffic, is explained in the *Multicast Routing in the VXLAN Underlay* section, *IP Fabric Underlay* chapter.

Layer-2 Multicast Forwarding

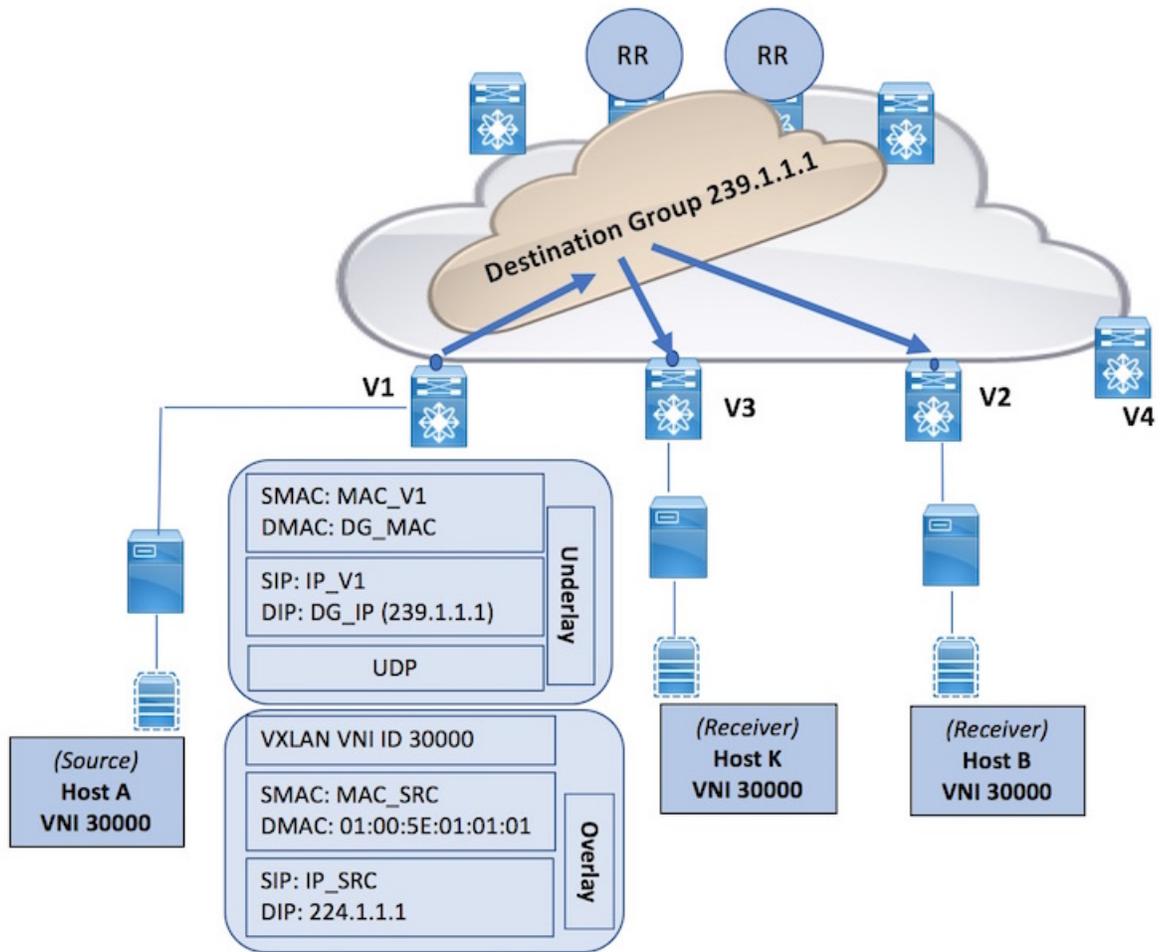
Some pointers are given below:

- *Multicast packets are sent only to participating leaf/ToR switch VTEPs*—A Layer-2 virtual network is tied to a unique multicast group. The L2 VNI-to-multicast-group configuration should be enabled only on those switches that have servers of this network connected. That way, only those leaf/ToR switches that have attached servers for the L2 VNI will participate in the multicast group. This will avoid unneeded traffic flow across the fabric.
- You should configure the same Layer-2 VNI to multicast group mapping on all VTEPs where the Layer-2 VNI has presence in.
- The multicast group is used for sending multi destination traffic, like broadcast (ARP requests), unknown unicast, and multicast traffic replication.



Note If Layer-2 multicast packets are to be sent to receivers in another data center, then a Layer-2 Datacenter Interconnect (DCI) technology like OTV is used to connect the two data centers and send the information.

Figure 1: Layer-2 Multicast Forwarding



- In the above example, let us say Host A, attached to V1, belongs to L2 virtual network 30000. With the below sample configuration, VNI 30000 is tied to the underlay multicast group or Destination Group 239.1.1.1.

(config) #

```
interface nve 1
  member vni 30000
  mcast-group 239.1.1.1
```

The Layer-2 virtual network has presence on V2 and V3, and the above configuration is applied similarly on V2 and V3 too.

- When A sends multicast traffic, V1 receives it and sends it to V2 and V3, and all other leaf/ToR switches that participate in VNI 30000.

Layer-2 multicast traffic flow example

Some details are given below for better understanding:

- A sends multicast traffic to its attached ToR/leaf switch V1. The DMAC and DIP of this packet/frame contains the server side multicast MAC and IP address (224.1.1.1).
- V1 knows that Host A, the source, belongs to L2VNI 30000. V1 VXLAN encapsulates the multicast packet and sends a single copy of the multicast packet upstream. Within the spine switch, the traffic to V2 and V3 is replicated efficiently (single packet per neighbor). This only happens as V2 and V3 are subscribed to the underlay multicast group (Destination Group) 239.1.1.1.

Note that, in a VXLAN multicast packet, the DMAC and DIP of the outer header contains the MAC and IP address of the destination multicast group

- Upon receiving, V2 and V3 decapsulate the VXLAN packets, check for intended multicast recipients, and send the original frame/packet to attached servers that are subscribed to the server side multicast group 224.1.1.1.



Note More than one Layer-2 VNI can be associated with the same underlay multicast group (Destination Group). In case multiple virtual networks (say, VNIs 30000 and 30001) are assigned to the same group 239.1.1.1, and there are VTEPs that serve only end hosts of 30000 and not 30001, those VTEPs will still receive multicast traffic for 30001's end hosts, but drop those packets if the Layer-2 segment is not instantiated. This is because all VTEPs that subscribe to a multicast group need to receive traffic for that group.

Multicast Communication in virtual Port Channel (vPC) Scenarios

A common, virtual IP address (VIP) is assigned to represent both the vPC peers for receiving and sending multicast (and unicast) packets.

- When a multicast source is behind a vPC complex, the VIP will be used as the outer source IP address of the VXLAN multicast packet.
- When a multicast receiver is behind a vPC complex, the receiver is announced by the VIP that represents the vPC pair. Multicast traffic intended for this receiver will only be sent to one of the two vPC nodes, since only one node will be active in the underlay multicast tree.

Layer-2 Multicast Configuration in the VXLAN BGP EVPN fabric

The following is a sample Layer-2 multicast configuration for two Layer-2 virtual networks with VNIs 30000 and 30001.

**Important**

This is a partial configuration in the overall VXLAN BGP EVPN fabric configuration, and should not be done in isolation. For complete configuration, refer the *Forwarding Configurations* chapter, *Cisco Nexus 5600 Series switch configuration* and *Cisco Nexus 7000 Series switch configuration* sections. These sections contain complete unicast and multicast configurations.

(config) #

```
feature pim
interface nve 1
  member vni 30000
  mcast-group 239.1.1.1
```

(config) #

```
interface nve 1
  member vni 30001
  mcast-group 239.1.1.2
```

**Note**

Ensure that you retain the same Layer-2 VNI to multicast group mapping on all VTEPs where the Layer-2 VNI has presence in.